

International

SRI
International

INTEGRATED SPLIT/MERGE IMAGE SEGMENTATION

Technical Note 441

July 1988

By: Kenneth I. Laws, Computer Scientist

Artificial Intelligence Center
Computer and Information Sciences Division

SRI Projects 2000 and 8388

The work reported herein was supported by the Defense Advanced Research Projects Agency under Contract Nos. MDA903-86-C-0084 and DACA76-85-C-0004.



333 Ravenswood Ave. • Menlo Park, CA 94025
(415) 326-6200 • TWX: 910-373-2046 • Telex: 334-486

Abstract

The KNIFE segmentation algorithm interleaves splitting and merging of regions during monochrome or multiband image partitioning. KNIFE splits regions along object boundaries, thus avoiding rectangular quadtree artifacts and establishing a context for good statistical decisions. Its iterative subregion extraction is based on multiband cluster analysis, with histogram-based threshold analysis used as a heuristic shortcut in simple cases. Splitting and merging decisions are based on sloped (rather than constant) surface fits, with successively more powerful thresholds and techniques employed until each region is split or found homogeneous. The user specifies only a desired level of segmentation, which is converted to procedural form by the KNIFE control process. The KNIFE package also offers a region-growing algorithm based on recursive splitting of neighboring regions. Examples of the two techniques are given for the domains of aerial cartography and reconnaissance, target cuing, and navigational vision.

Contents

1	Introduction	1
1.1	Recursive Partitioning	1
1.2	Recursive Merging	3
2	KNIFE Segmentation Techniques	5
2.1	Data Transformation	5
2.2	Histogram Analysis	7
2.3	Median Splitting	10
2.4	Cluster Analysis	11
2.5	Spatial Analysis	13
2.6	Noise Cleaning	14
3	KNIFE Region-Growing Techniques	16
4	Future Research	18
5	Examples	19
6	Summary	31

1 Introduction

Pixel grouping (or *feature extraction*) is a useful first step in digital image analysis. Vision scientists have developed many algorithms for image segmentation, but no one approach is adequate in all cases. My goal is to combine the best feature extraction operators in a single system, with intelligent control processes to exploit their individual strengths and compensate for their weaknesses.

The KNIFE analysis system developed at SRI International integrates region splitting and region growing in a successive-refinement partitioning system, with multiband cluster analysis used to guide the partitioning steps. Each component algorithm is newly developed and has advantages over previously documented approaches, but it is the integration of the techniques that constitutes notable progress in digital image analysis.

The KNIFE segmenter takes either monochrome or multiband image data, with bands encoding either spectral or textural information. An early version has been used as a preprocessing stage for shape-based recognition of buildings, roads, and trees in aerial images [Fua 86, 87ab]. Other applications, such as tracking of objects in ground-level sequences [Laws 88b], are under development. The software system includes many display and analysis tools (including a new "coarse coding" statistical classification technique [Laws 88a]). This paper describes only its integrated split/merge capabilities.

1.1 Recursive Partitioning

Simple objects, such as printed characters, may be recognized by means of correlation-based template matching. Complex scenes with unpredictable object configurations require more sophisticated analysis. Image pixels must first be grouped into lines and regions so that high-level techniques have meaningful "chunks" to reason about. These extracted features should be few in number but high in information content.

Threshold segmentation is one of the oldest and most successful techniques for feature extraction. Isolated objects that are consistently darker or brighter than the rest of a scene (such as stained cell nuclei or infrared hot spots) can be found by simple brightness thresholding [Prewitt 66, 70; Nakagawa 78b; Weszka 78ab; Otsu 79], with histogram analysis typically used to select the threshold.

A less obvious target—of intermediate brightness, but contrasting with its immediate backgrounds—can also be extracted by global thresholding, say, by trying different thresholds until a resultant region is recognized as a scene object [Selfridge 82]. The threshold that separates such a subregion from its background may cause other subregions to break into meaningless pieces, so the system must recognize and deal with the two cases. Parallel techniques are required if suitable thresholds are to be chosen for all scene objects in reasonable time.

Sometimes an extended object, such as a river or road, is brighter than one part of its background and darker than another part. This can also be true of replicated objects (barracks, suburban houses, trees, corn rows, etc.) from a single target class. Histogram analysis may reveal good pairs of thresholds if the scene is simple, but must be combined with adaptive or recursive thresholding for complex scenes [Prewitt 70; Ohlander 75, 78; Price 76; Nakagawa 78a].

Adaptive thresholding varies from one part of an image to another, and so can extract sloped or slowly changing regions—at least in theory. Effective threshold functions, or

surfaces, are difficult to construct. Edge-based threshold selection [Panda 78; Milgram 79, 81; Weszka 79; Kohler 81; Minor 81; Hartley 82] techniques are dependent on the gradient operators used, and no one edge filter is likely to work everywhere in a scene. Sequential boundary tracking [Montanari 71; Pingle 71; Shirai 72; Martelli 72, 73, 76; Ramer 75; Nahi 77, 78; Grinaker 80] seems necessary for adaptive partitioning, but is inherently slow and requires high-level models of the structures being sought. Edge-element linking [O'Gorman 73; Nevatia 76, 80; Wechsler 80; Walters 86] is currently favored over adaptive thresholding, despite its inability to find meaningful closed regions in most natural imagery. Either specific [Bolles 78, 81, 82, 86; Levine 81a; Binford 82] or generic [Kelly 70; Bajcsy 74; Sakai 76; Sloan 77, 79; Quam 78; Agin 79; Brooks 79, 81, 84; Russell 79; Fischler 81a, 82; Mackworth 81; Selfridge 81, 82; Glicksman 83; Havens 83; Harlow 84; Khan 84; Nagao 84; McKeown 85; Lowe 85; Fua 86, 87ab] object models are needed to complete the extraction of scene objects.

Edge-based methods exploit local discontinuities, but ignore information in the interior surfaces of regions. Area-based splitting permits subregions to be detected even when edge detectors find no boundary between them [Silverman 86, 87]. A good adaptive thresholding system would combine the positional accuracy of edge linking with the robustness of statistical surface fitting, but there has been little progress in this direction. My own research in area-based segmentation uses repeated splitting and merging to refine region boundaries, but does not yet incorporate edge linking or line tracking to hypothesize partitionings. (See Fua and Hanson [Fua 86, 87ab] for one such approach.)

Recursive thresholding exploits region statistics to control partitioning. The entire image is first split, then each extracted region is considered for further splitting. Each successive split improves the context in which additional decisions are made. Segmentation of any scene area can be terminated when its regions are known to a sufficient level of detail. Recursive thresholding has been quite popular for partitioning images into compact, homogeneous regions [Yachida 71; Robertson 73ab; Ohlander 75, 78; Price 76; Riddler 78; Shafer 80, 82; Laws 82], but does tend to fragment such slowly varying regions as river and road networks.

Recursive partitioning is a more general term for the repeated splitting of regions. It includes recursive thresholding, but may also exploit cluster analysis, pixel classification, linear feature extraction, and template or model matching to extract subregions [Tomita 73; Tsuji 73; Zucker 75]. A recursive segmenter typically produces a tree of sub-region relationships, although such trees are of little use. (The order in which data bands and thresholds are selected for partitioning is unstable. The important question is not how a region was discovered, but why it should be retained.)

Recursive partitioning (or splitting) is particularly useful because it generates a series of meaningful intermediates as it works toward a full scene parse. A well-designed segmenter will first locate coarse partitions (e.g., sky, land, water), then refine them. The control process can restrict the level of detail sought in each region, using discovered context as a guide to further analysis. A "pyramid" of parses at differing resolutions may be saved for use by other analysis programs.

A disadvantage of recursive splitting is that it may take a long time to reach the level of small targets in a large image. Targets with distinctive spectral values can be located in one classification or thresholding step, but objects similar to the background statistics must be found by successively paring away other regions. This sequential partitioning is

slow unless control processes can focus quickly on the image areas containing important information. Top-down segmenters without semantic guidance should be used for target detection only if a full scene parse is also necessary for other purposes.

Basic operations in recursive splitting are tentative partitioning of a region into two or more subregions and testing to determine whether to accept the split. Blind search through the space of all possible segmentations is infeasible, so we require a method of generating likely subregion boundaries. A reliable hypothesis generator permits us to accept or reject a single "best" partitioning instead of searching more deeply through multiple sequences of operations. Forward search, or hillclimbing, is then sufficient, with later splitting and merging taking the place of backtracking to escape from errors.

Any number of segmentation approaches can be used to generate tentative subregions for recursive partitioning. Each technique should complement rather than duplicate the test used to accept or reject each split. The KNIFE system uses a spectral splitting method combined with spatial validation of hypothesized splits. It is a purely syntactic, or low-level, segmenter that has no knowledge of target characteristics beyond the desired level of detail set by the user. (Control can be passed to a higher-level program or to an interactive user after each split/merge step, but the standard procedure is to segment autonomously until a specified level of detail is reached. This may be done over the entire image or within any initial set of subregions.)

Spectral methods search for reasonable clusterings of the pixel brightnesses (or multi-band intensities) within a specified region. It is presumed that adjacent objects will differ significantly in at least one pixel property, and that most spectral clusters correspond to spatially coherent subregions arising from significant scene structures. These heuristic assumptions are adequate in practice, although it is easy to show that the human visual system uses more sophisticated techniques.

The KNIFE segmentation algorithm can be limited to strict hierarchical segmentation if desired, but gains much of its power from interleaved region-merging steps that correct for the oversegmentation typical of thresholding and cluster-based approaches. KNIFE thus combines recursive splitting and recursive merging in a region-based iterative or *relaxation* partitioning system. Its built-in control processes bias the relaxation toward splitting or merging, with the complementary process correcting any obvious errors. An annealing schedule guarantees termination by making splitting more difficult in successive passes through the image.

1.2 Recursive Merging

Region growing and merging are techniques for expanding one set of regions at the expense of another. Region growing appends individual pixels (or sometimes groups of pixels) to specified *seed regions*, whereas region merging absorbs entire regions at each step. Seed regions that touch are generally not permitted to take pixels from one another, although relaxation-style algorithms can be formulated to allow this without risking infinite cycles.

Region growing is typically limited to one seed region, or to a few seeds identified as region centers or known material types. Although useful for extracting extended, slowly varying objects, such as rivers and roads (as well as objects seen in strong perspective), region growing suffers from lack of a reliable stopping criterion. It is really only useful

when exactly one seed region has been identified for each scene object, but that requires very strong object models. Even then, roads and other thin regions can be broken as other seeds grow across them—especially if the image is slightly blurred.

Region merging is identical except that all participating regions are active seeds and entire regions are absorbed at each step. This is typically used for partition editing, although one can do a full segmentation by starting with each pixel as a separate region. Merges may be tested in a particular scan sequence, queued according to a heuristic-quality metric, or accomplished gradually by iterative updating of connection probabilities.

An advantage of merging over single-pixel growth is that more context can be used in making each merge decision. Region statistics for both regions are available, as well as measures of boundary shape and strength. Maintaining these statistics and descriptors—along with neighbor lists, candidate queues, and semantic region labels—during merging can be difficult, although many clever representations and control strategies have been developed [Fennema 69; Brice 70; Jarvis 73, 77; Yakimovsky 73ab, 76; Gupta 74; Triendl 74; Hanson 75; Freuder 76ab, 77; Levine 76; Somerville 76; Zucker 76; Riseman 77; Raafat 80, 86; Asano 81; Pong 81, 84; Beaulieu 83; Suk 83; Silverman 86, 87].

Region growing and merging usually produce inconsistent levels of detail. Perceptually salient or semantically meaningful regions will merge with one another or with their backgrounds before other such regions have coalesced. Even if this were not the case, it would be difficult to determine when merging should cease. An obvious solution is to incorporate object recognition and thus semantic control. This has been tried for both region growing and region splitting [Yakimovsky 73ab, 74; Feldman 74; Tenenbaum 76, 77; Sakai 76; Nagao 77, 79, 80; Kestner 80, 82; Ohta 80; Levine 81ab; Sze 82], but current approaches have so far been inadequate to represent and exploit the predictabilities of complex natural imagery.

An alternative, used in KNIFE, is to combine region splitting with region merging in a competitive/cooperative relaxation paradigm. Each technique compensates for weaknesses of the other. Region splitting tends to oversplit (if all important regions are extracted); region growing tends to overmerge. When yoked together, they produce acceptable segmentations with predictable levels of detail. (Integration with pixel classification, edge-based analysis, and higher-level control would improve results even further, but is beyond the scope of this report.)

The KNIFE segmentation algorithm incorporates region merging to clean up fragmented “noise regions” created during splitting. Postediting of this sort has often been proposed and, in fact, is a standard feature of quadtree split/merge systems [Pavlidis 72, 75; Horowitz 74, 76; Feng 75; Tanimoto 77; PCChen 79, 80, 83; Grinaker 80; Jayasimha 81; Browning 82; Conners 82, 84; Pietikäinen 82b; Sze 82; Lee 83; Doherty 86; Bhanu 87b]. Unfortunately, merging with arbitrary neighbors destroys such hierarchical representations as quadtrees, making further splitting inconvenient. Continued splitting may require a separate quadtree for each region.

I have developed efficient data structures and algorithms for integrated multiband splitting and merging, as well as libraries of subroutines for manipulating data bands and image windows, region maps, multispectral histograms and clusters, line segments, points, lists, heaps, and related entities. These routines constitute a spatial-representation language used in KNIFE to maintain an acyclic directed *region graph*¹ describing the

¹The partitioning algorithm produces a hierarchical tree structure, but the user or supervisory process

current state of a segmentation. Tentative splits or merges are considered until one is successful enough to be incorporated in the region map and its graph; then the process recommences within each new or altered subregion.

The KNIFE system includes a top-level region-growing command separate from the segmentation noise-merging process. This algorithm also integrates splitting and merging, but with merging as the dominant process. I shall first describe the segmentation algorithm and its use of merging, then the region-growing algorithm and its use of segmentation.

2 KNIFE Segmentation Techniques

The KNIFE segmentation algorithm consists of six techniques: data transformation, histogram analysis, median splitting, cluster analysis, spatial analysis, and noise cleaning. Data transformation is the only step currently requiring human expertise, although the KNIFE system works well enough on even monochrome imagery to make data transformation usually unnecessary. The other techniques are fully automated and depend on only a single *segmentation-level* parameter (or user-specified level of detail). The techniques themselves are rather complex and are invoked in data-dependent sequences. I shall try to explain what they do—and why—without getting down to the level of flowcharts or program code.

2.1 Data Transformation

A problem with recursive splitting is that the initial partitioning of a whole image can be very difficult. I call this the *dead center* problem, since getting the segmenter started is much like starting a wheel with its crankpin on dead center. The histograms of a large and complex image often look Gaussian, defeating simple thresholding schemes. Many researchers have used color or texture bands to increase the information available for region splitting [Haralick 69, 75; Yachida 71; Nagy 72; Robertson 73ab; Schachter 75; Coleman 77, 79; Goldberg 77, 78; Narendra 77; Yoo 78; Bryant 79; Fukada 80b; Matsumoto 81], although this is often insufficient for scenes with many objects. (Robertson also split images along grid lines, segmented, and then sewed the quadrants back together in a manner foreshadowing quadtree segmenters.)

Ohlander used redundant color transforms to search for any isolated histogram peak. A peak that is obscured in one marginal histogram may be completely isolated in another. Price added a planning system [Kelly 70; Hanson 75; Price 76; Nagin 77; Reynolds 84] to this approach, reducing processing time to a tenth by projecting a low-resolution segmentation into the full image so as to limit areas over which histograms were computed. A modern equivalent might be to filter the image, then extract areas within zero crossings or with suitable Gaussian curvature as tentative regions [Seibert 88].

The KNIFE splitting algorithm is based on univariate histogram analysis and multivariate cluster analysis. Easy splitting decisions are made by thresholding at single-band histogram valleys. When all such attempts fail, the program thresholds each data band at its median and selects the best result. (This may fragment regions that span the arbitrary threshold, but further splitting and merging can produce an acceptable partitioning.) If no

can introduce arbitrary links to indicate semantic relationships. The software therefore allows for regions to have multiple parents.

initial split satisfies the built-in statistical criteria, KNIFE resorts to full multidimensional cluster analysis.

Individual color bands are *marginal distributions* from a three-dimensional color space. (Scene objects may have infinitely varied spectra, of course, but human color vision is limited to a three-parameter summary. Our color cameras and other systems are engineered to duplicate this perceptual bias, and typical evolved or manufactured objects are likewise matched to the sensitivities of biological vision.) For multidimensional cluster analysis, it should make little difference which coordinates are used to represent pixel colors.

In practice, however, there are advantages in dealing with statistically independent color bands. It is difficult to explain orthogonality or independence of color parameters in a rigorous manner. Each band should convey independent information about scene objects, but the bands are necessarily correlated for pixels within each object. One view is that bands should match our perceptual channels, with changes in any one parameter value not affecting our perception of the others. Another view is that bands should approximate the principal components, or eigenvectors, of the color space—relative to the color distribution within a single image or an ensemble of images. Everyone would agree, though, that HSI (hue, saturation, intensity) bands are more independent than RGB (red, green, blue) bands for typical imagery.

The KNIFE approach of using single bands to make obvious decisions works best with independent bands. The current cluster-analysis algorithm also expects independent bands, although it can function with ordinary RGB input. The clustering technique could be made more sophisticated, but it is simpler to transform RGB input to a better representation. I use a new VHS (vividness, hue, saturation) transform [Laws 88b]. Vividness is similar to intensity or brightness, but gives primary colors the same value as a pure white. Hue is the same as for the HSI (or IHS) system, except that the origin is rotated to the purple region. Saturation differs from the formula commonly used in computer vision in that colors near pure black and pure white are assigned low saturation. The VHS coordinates are not necessarily optimal—blue should not be as vivid as yellow, for instance—but they are easy to compute and work very well for segmentation.

For monochrome imagery, pixel-based cluster analysis is very similar to histogram analysis. Some cluster algorithms are influenced by neighboring pixel values [Narayanan 81; Bhanu 82; Davis 82], but KNIFE's approach gives much the same results as its heuristic valley finder—except for the rare cases when segmentation can proceed only by using two thresholds simultaneously. KNIFE therefore skips the clustering step for monochrome imagery unless the user specifies otherwise.

If monochrome imagery must be segmented, however, and if histogram analysis is too weak, there are still ways of constructing additional data bands. The point-by-point logarithm of weighted local variance, for instance, is a powerful measure of image texture. (This and all texture measures used by me are computed with a binomial or Gaussian weighting that emphasizes the center of each data window.) When gray level and texture level together are insufficient, more complex texture bands can be used. The SRI IU Testbed includes an improved version of my texture energy measures [Laws 80ab, 88a] with binomial weighting to reduce rectilinear artifacts. Gabor filters and fractal texture operators are also available. Other texture measures may be important for recognition and higher-level analysis, but it would be very unusual to find adjacent objects differing in high-order texture and not in brightness or local variance.

Special-purpose object detectors (e.g., for roads) can also be used to generate data bands where each pixel is marked with the probability of that object's being present. The KNIFE package does not include such detectors, but can create signature similarity bands [Laws 85, 88b] from user-supplied training examples. Semantic "treeness," "grassness," or "houseness" scores, for instance, can be computed from a monochrome or multiband image and a set of prototypical regions. The KNIFE splitting algorithm can then utilize the semantic bands to facilitate segmentation.

There is one more possibility I should mention. It is easy to compute local gradient direction and magnitude from an image for any specified neighborhood size and weighting. The local variances of these quantities can be used as texture measures. (An information-theoretic entropy measure can be substituted for a circular measure, such as gradient angle.) It would seem that the gradient measures themselves could also be employed as data bands, suitable for extracting image regions of constant slope.

This approach may have some merit for robotic vision and industrial inspection, although surfaces and shadows will typically exhibit spatial nonlinearities in their brightness patterns. My experience has been that gradient measures over typical aerial imagery are too noisy to be useful. While they may help extract a few objects, they interfere with extraction of many others. Gradient bands require special techniques (as do range bands, radar data, and various other forms of imagery). I have built knowledge of region intensity slopes into KNIFE's merging algorithm; this seems to work better than providing pixel-level gradient bands to the splitting algorithm.

2.2 Histogram Analysis

One way to partition a region is to divide the pixels into those above a threshold and those below it, as discussed above. Modern computers are fast enough to try all possible single-band thresholds, but connected-component extraction and spatial analysis still take considerable time. It is best to examine the pixel statistics and test only the most promising thresholds.

I shall assume for the moment that each scene region² contains pixels with a Gaussian distribution of intensities, or perhaps a sum of Gaussian distributions. (An image of a field, for instance, might include interspersed populations of ground and vegetation pixels.) Each image region should therefore contain a sum of such Gaussians. The job of histogram analysis is to discover brightness thresholds separating one or more source populations from the rest. Some errors that are due to later split/merge editing can be tolerated, but it is best to choose thresholds in a sequence that minimizes overall boundary misplacement.

Let us suppose that the single-band standard deviations of the peaks (caused by either texture and noise or by a slow variation in brightness) are small compared with the separation between peaks. It is then easy to locate the histogram valleys that correspond to reasonable thresholds. One good rule is to place thresholds where the smoothed histogram has zero first and third derivatives and positive second derivative, for some maximal amount of smoothing. Many other heuristics exist for partitioning histograms, and all work well in this simple case.

²Scene regions are semantic units that are independent of the extracted image regions.

A complication occurs when a peak for one scene region falls between two peaks for interspersed pixel populations in a second scene region. Sequential application of two thresholds will then fragment the bimodal scene region. The Ohlander and Price solution is to apply two thresholds to the image simultaneously, extracting only the regions between them. An early version of KNIFE permitted such multiple thresholds, but the control process was not able to recognize when they were appropriate. Applying any one threshold introduces noise into a complex-image partitioning action; multiple thresholds tend to introduce far more noise. My current system depends on texture or signature-similarity transforms, cluster analysis, statistical goodness-of-fit tests, and *delayed commitment* to handle this rare situation.

Monochrome cluster analysis, as already mentioned, can sometimes solve the problem (if the central peak is large enough), but is seldom worth the expense. Delayed commitment, or procrastination, means rejecting unsatisfactory partitionings and trying again after other operations have changed the context of the decision. Region splits due to other histogram valleys, or merges due to splitting of neighboring regions, may alter the histograms enough to allow a reliable threshold or cluster to be found.

Someday the control system may be sophisticated enough to use signatures of objects found elsewhere in an image to control the extraction of difficult regions [Narayanan 81; Conners 82, 84; Trivedi 84ab; Harlow 85; Weymouth 83; Laws 88b]. At some point, however, any control process must conclude that a region is spatially homogeneous even if the histograms shows spectral subpopulations. The current KNIFE system does a good job of terminating splitting attempts that are likely to be futile.

A more common problem arises when subpopulation standard deviations are large and histogram peaks overlap. Any threshold will then cause some pixels to be cut off from their source populations. Such pixels, lying on the outer borders of the original region, or inside the newly formed major subregions, will be isolated as little noise regions. Those lying along the new split will be included with the wrong major subregion, but may form isolated noise regions during later splitting operations if a sufficiently high level of detail is sought. A merging or editing step is needed to identify these noise regions and merge them with appropriate neighbors [Tanimoto 76, 77; Lumia 81, 83].

Many methods, including all those employed for edge-based partitioning,, have been developed to combat this fragmentation. One approach, called conservative cluster formation [Nagin 77], uses a pair of closely spaced thresholds bracketing each histogram valley. Gray levels within this bracket are pulled out as separate noise regions, with further processing required to split the border regions and merge the resultant pieces with appropriate neighbors. The KNIFE program offers an optional golden-section search to optimize its heuristically chosen thresholds. (The "optimization," based on minimizing the total noise-region area, yields poor results.) One can also search for thresholds that will segment along image curves of highest gradient [Milgram 79, 81; Barrett 81; Hartley 82], although any such threshold tweaking will be inferior to local adjustment of border positions by gradient hill climbing.

Yet another cure is relaxation enhancement, with image pixels iteratively adjusted to have gray levels closer to those of their spatial neighbors [Tomita 77; Nagin 78, 79, 82; Rosenfeld 78, 79a; Danker 81; Narayanan 81; Bhanu 82, 87b; Davis 82; Hartley 82; Laws 83; Parvin 83]. This tends to flatten regions and sharpen their borders, simplifying the partitioning task. (Median filtering, or any variant of edge-preserving smoothing, can

sometimes produce a similar effect at less cost.)

Relaxation methods use a pixel's local context to determine which side of a threshold it belongs on. Texture operators similarly gather local context into vectors of pixel descriptors to be used in thresholding or clustering. Texture bands, which resemble blurred images, have been used in the same manner as spectral descriptors for segmentation and labeling [Tomita 73; Tsuji 73; Thompson 74; Triendl 74; Hanson 75; Pavlidis 75; Zucker 75; Carlton 77; Coleman 77, 79; Keng 77; Deguchi 78; Mitchell 78, 79; Tsai 78; PCChen 79, 80, 83; Harlow 79ab, 85; Rosenfeld 79b; Schachter 79; Wechsler 79, 80; Grinaker 80; Laws 80ab, 85, 88a; Raafat 80, 86; Burt 81; Jayasimha 81; Pietikäinen 81, 82a; Connors 82, 84; Davis 82; Lee 83; Trivedi 84abc, 85; Silverman 86, 87]. (Other researchers have used segmentation as a preprocessing step for texture analysis [Abele 80; Lumia 81, 83].)

Multivariate histogram or cluster analysis is powerful, but segmenting even a bivariate histogram is almost as difficult as segmenting an image. Multivariate histograms are also too large to be stored conveniently. The KNIFE program uses multivariate clustering as a final test of homogeneity for each region, but performs the bulk of its analysis on univariate histograms. The univariate histograms of all but very small regions are saved, or cached, so that they can be used during merging and other operations. This saves considerable time.

The error correction inherent in KNIFE's integrated split/merge approach yields single-band segmentations as good as those achieved by Ohlander with three original and six derived color bands. Additional color, texture, or signature-similarity bands provided by the user can help with difficult discriminations in any task domain. Neighboring image areas that differ significantly in even a single measurable property are considered separate regions arising from different generating processes.

Many methods exist for partitioning univariate histograms. Techniques for decomposing arbitrary mixture densities into true Gaussian components are unnecessarily complex, especially since the underlying assumptions are violated in most real images. Even if scene objects were homogeneous and if highlighting, perspective, sensor saturation, and picket-fence effects (due to contrast stretching or color transformation after digitization) could be ignored, signatures of regions derived by repeated thresholding would not be Gaussian. Such histograms tend to be uniform within narrow intervals, with outlying spikes resulting from merged noise regions. Very small regions often have exceedingly noisy histograms that defy even human-guided analysis. Fortunately, there is no need for an underlying model of histogram formation. All that is necessary is a way of selecting the most obvious histogram valleys first, with less obvious thresholds suggested if necessary. This can be done by a scale-space analysis [Witkin 83; Carlotto 85] or by successively reducing histogram smoothing while searching for useful peaks or valleys.

The KNIFE system uses a modified version of the PHOENIX heuristics [Shafer 80, 82; Laws 82] for identifying valleys, gradually reducing histogram smoothing and relaxing threshold acceptance criteria until a good threshold is found in at least one image band. (Almost any one-dimensional segmentation or clustering procedure could be used.) The PHOENIX heuristics work well, furnishing good candidates for the spatial analysis used to verify each split. If more than one acceptable threshold is found, the one giving the best spatial goodness-of-fit score is used.

The following takes place at each stage, or *segmentation level*, until the region is split

or failure has occurred at a user-specified maximum level: Each single-band histogram is smoothed and all resultant local minima are marked as potential thresholds. The intervals between these minima are examined; any interval that does not pass certain tests (e.g., adequate included area and peak-to-valley ratio) is merged with its neighboring interval.³ on the side with the higher valley. This is repeated until all the remaining intervals satisfy all of the screening tests. Any intervals that survive are then passed to a spatial analysis routine, described below, that thresholds the image region and computes a goodness-of-fit score. The best threshold across all data bands is chosen. If its score is high enough, the partitioning is retained. Otherwise the cycle repeats at the next segmentation level using reduced smoothing, relaxed screening criteria, and a lower permitted goodness-of-fit score. Only when all attempts have failed does the histogram analysis system declare the region homogeneous in all bands and pass it on to more powerful procedures.

This segmentation procedure, together with its data structures and interfaces, took several years to develop and automate. The dependence on heuristic thresholds and local decisions (in lieu of global optimization) may bother purists, but the search mechanism works well and seldom selects a poor threshold. Search techniques are at the heart of most artificial intelligence programs, and seem appropriate and effective here. If parallel processors are being employed or the human vision system modeled, however, I would do simultaneous analyses at all segmentation levels and then combine them or select the best [Fua 86, 87ab].

2.3 Median Splitting

Histogram analysis may fail for two reasons: the region may be truly homogeneous, or it may be too complex to have separable histogram peaks in any one data band. The first case can be confirmed only by trying additional splitting methods, such as multivariate cluster analysis, until all reasonable partitioning techniques have failed. The second case will also yield to this approach, but I have found that a quick heuristic test, a *median split*, is often just as effective.

Complex regions, such as entire aerial images, often have histograms approximating broad Gaussian distributions. This is due to the *central limit* effect, whereby sums of many independent distributions (arising from numerous scene regions) tend to be Gaussian. Even multivariate spectral analysis might have trouble partitioning such distributions.

We could invoke a nonspectral technique, such as region growing or boundary following, to get the segmenter off dead center. At present, however, the KNIFE segmenter simply tries the region's median (in each data band) as a threshold. Most of the subregions in a complex scene will fall above or below this level, so will be extracted correctly. Large or extended regions that include the median intensity will usually split into coherent pieces that are easy to remerge in a later step. Large textured subregions may fragment, but the goodness-of-fit score will then prevent the split from being accepted. (A different band, such as an appropriate texture transformation, may succeed in extracting the textured subregions.)

³Certain data dimensions, such as hue, should be considered circular. This is not difficult, but has never been implemented in KNIFE. The system relies instead on color transformations that seldom generate gray levels near the ends of the histogram. Gradient directions cannot be handled so easily, but seem not to be a useful transformation for segmenting complex natural imagery.

The only real danger in median splitting is that important small subregions, either textured or slowly varying, will be fragmented. The goodness-of-fit score is not affected by the fate of small subregions if larger or more numerous ones are well fit. This is a danger with any threshold split, although the risk is slightly higher when the threshold is at a histogram peak instead of a valley. Fortunately, these situations seldom arise. With luck, the KNIFE algorithm can recover later from errors by continued splitting and noise-region merging.

In short, median splitting is a quick and useful test for region substructure. The few errors it makes are usually corrected by subsequent splitting and merging operations. Its speed can be exploited precisely because the KNIFE system integrates splitting and merging.

2.4 Cluster Analysis

A single data band containing distinguishable pixel distributions is unavailable often enough to interfere with recursive splitting. We would like to detect objects with unique spectral signatures regardless of any overlapping background signatures. This requires pixel classification based on strong object models or full cluster analysis. I have taken the latter approach, as have many other researchers [Haralick 69, 75; Nagy 72; Triendl 74; Hanson 75; Schachter 75, 77, 79; Carlton 77; Coleman 77, 79; Goldberg 77, 78; Nagin 77, 78, 79, 82; Narendra 77; Riseman 77; Saba 78; Yoo 78; Bryant 79; Mitchell 79; Wechsler 79, 80; Abele 80; Fukada 80ab; Raafat 80; Rassbach 80; Jayasimha 81; Lumia 81, 83; Matsumoto 81; Sarabi 81; Davis 82; Wharton 83; Trivedi 84c, 85].

Cluster analysis is often used to select region centers that are then grown. In other cases, tentative regions found by oversegmenting are clustered in a spectral [or measured-feature] space to find a final set of regions. I use clustering only for splitting regions, with connected-component analysis, region merging, and perhaps further splitting used to verify and improve the segmentation.

Ohlander, Price, and others have depended on redundant color transformations to supply the spectral separability required for good segmentation. The intuitive conclusion, which can be traced back at least as far as Yachida and Tsuji [Yachida 71], is that overlapping peaks in the original single-band histograms may be well separated in other views of the same multidimensional space. Ohlander used RGB, HSI, and YIQ color bands together to achieve sufficient separability of peaks in the three-dimensional color space. This combination has continued to be popular despite severe problems with transformation singularities, dynamic range (especially the Q band), and picket-fence notching [Kender 76, 77].

Analysis of redundant data bands is a heuristic approximation to full multivariate cluster analysis. Ohlander found it faster to search for peaks in nine histograms and to apply a resultant pair of thresholds than to perform a three-dimensional cluster analysis and evaluate a discriminant function at each pixel. I have found a faster and more robust method: using individual data bands for most region splits, but invoking full cluster analysis whenever histogram analyses (and median splits) fail. It is faster—on sequential computers—because fewer data bands need be manipulated, and because cluster analysis and integrated merging result in cleaner regions that require fewer splits.

For clustering I use a modification of the ISODATA algorithm [Ball 65, 67]. My version,

ANISODATA, uses cluster kurtosis (or fourth moment) in each spectral band to control cluster splitting. It is designed for anisotropic data spaces—such as VHS, HSI, or YIQ color coordinates—where each data dimension has a different variance. (An alternative, of course, is to normalize each data band to a standard variance. This requires floating-point image representation and increases the data-handling complexity of a multiband analysis.) ANISODATA makes no special provision for correlated data bands—such as RGB coordinates—where information about one band value implies information about another.

I start by choosing cluster seeds at the data centroid and at the centroid plus or minus three standard deviations in each dimension. These statistically defined points are easily computable from the stored region histograms, but there are other ways of selecting seeds (e.g., randomly sampled or spatially distributed pixels [Yakimovsky 73ab]; sequential selection of pixels that are unlike their previous counterparts; recursive use of a previous cluster segmentation [Hanson 75; Trivedi 84c, 85]; or use of subregions found by another method [Haralick 69, 75; Beaulieu 83]).

Up to four passes through the region data are permitted. All pixels in the region are first assigned to cluster centroids. Very small clusters are eliminated and the remainder then subjected to the usual ISODATA splitting or merging. (Here too variations are possible. A sample of pixels could be used to save time, particularly during the early passes. Clustering of subregion descriptors will always be faster than clustering of pixels, but requires fast methods of oversegmenting and of computing descriptors.)

The ISODATA split/merge schedule is complex and differs from one pass to another. In my version, clusters with kurtosis of less than 2.5—where a Gaussian is 3.0—in any dimension are split, unless there are too many clusters already. (The target number is two, but the algorithm can produce more.) If no splitting has occurred, clusters within three standard deviations of each other may be merged. (Multidimensional Mahalanobis distances from each cluster are computed, with merging permitted if the smaller distance is less than 3.0. Merges are performed best-first, with no cluster merging more than once, until at most four merges have occurred.) These parameters bias the analysis toward splitting; later spatial verification and merging will prevent serious errors.

KNIFE then combines all the clusters except the largest. Connected components are extracted and the usual noise cleaning takes place, controlled by the user-specified level of segmentation detail. Each of the resultant subregions becomes available for recursive segmentation. Lumping of all smaller clusters sometimes sacrifices processing time, but pulling out one pixel population at a time in this manner usually results in the best segmentation.

Segmentation of a region terminates when all splitting methods fail (at a specified segmentation level). It follows that every pixel in the image must ultimately be subjected to full cluster analysis. Although this may seem like a great deal of computation, it should be noted that the cluster analyses are done region by region—usually after other methods have reduced each region to a small number of source populations. Clustering proceeds very rapidly under such conditions and the overhead of this final check is not high. It is true that full cluster analysis is sometimes required to effect the initial split on a large and complex image and that this can take a great deal of time, but it is better than being unable to split an image at all.

2.5 Spatial Analysis

Acceptance of a partitioning, whether from thresholding or from cluster assignment, depends on a spatial quality score. Desirable splits usually produce several nearly equal subregions, plus a fair number of noise fragments. Not all small subregions are noise, however; even a single pixel that differs strongly from its surrounding distribution may be semantically important. Developing a metric that reflects these facts was difficult, but a simple statistical approach has now been found to work well.

The first step must be the extraction and representation of each subregion. Regions could be represented by boundary chains, quadtrees, or other data structures. KNIFE uses label maps as primary representations, with bounding rectangles and optional histograms, subregion lists, and other descriptors attached to a *region record*. The problem is this: given an image region in which each pixel has one label from a set, a map of maximal connected subregions and a database of subregion descriptors must be produced.

Early versions of KNIFE—then called SLICE—used a modification of the connected-component extraction routine from the PHOENIX segmenter [Shafer 80, 82; Laws 82]. This proved too slow despite careful optimization of the computer's memory allocator. It also consumed too much active storage when several alternative partitions with thousands of subregions had to be retained, and resulted in further wasted processing time because of the fragmentation of virtual memory. I have now developed a much faster algorithm.

Connected-component extraction, also known as region coloring, requires at least two passes through the thresholded image or cluster map—the first to tentatively label pixels and determine label equivalences, the second to assign final subregion labels. (Sequential boundary tracing would require even more computation.) KNIFE computes subregion statistics during the first pass, aborting the analysis if results are unsatisfactory. It also discards all but the best of the single-band parses before starting its final pass and instantiating the subregions as new image regions. This saves considerable time over repeated full extraction, at least on sequential computers. It also simplifies the algorithm and avoids having to create and manipulate competing subregion maps and descriptors.

Subregion descriptors for efficient parse pruning are computed from subregion fragments, or *patches*, as they are encountered during the first pass. A tree of patch records is maintained and is collapsed to a tree of subregion descriptors at the end of the pass. These descriptors can include subregion area, minimum and maximum coordinates, shape moments, and gray-level statistics. Subregion adjacencies could also be tracked, but this is not being done at present.

KNIFE currently computes subregion area and surface-fit coefficients. Residual variance after linear fit to each subregion is then computed and compared with that for the original region. An approximate F statistic represents the improvement in surface fit that is due to the freedom of fitting separate coefficients for each subregion. The F ratio is adjusted for the total number of subregions, thus penalizing low-variance fits caused by excessive fragmentation. This metric performs better than previous ones based on just the distribution of subregion sizes. It also requires only one parameter, a critical significance level, rather than the task-dependent minimum and maximum "target size" parameters required in earlier program versions.

The test metric has the form of an F statistic, but does not have an F distribution. This is because the hypothesized subregions are not randomly selected, but are constructed for

maximum discriminability. No theoretical derivation of the true distribution is possible, but extensive experimentation has yielded critical values that work well for any desired level of detail. Hypothesized partitions with larger values are accepted; those with smaller values are rejected.

If all threshold and cluster partitions are rejected, a region is presumed homogeneous and is marked with the current segmentation level. This stored *seglevel* value blocks useless attempts at resplitting during later operations, but permits splitting at finer levels of detail. The stored value is reset to zero if any neighboring region is ever merged with this one, even if the neighbor is very small.

2.6 Noise Cleaning

If a threshold or cluster split is accepted, each pixel's label in the region map is set to that of its new subregion. Neighbor relations are determined (by means of a four-connected grid) and additional subregion descriptors, such as histograms, may be computed. The subregions are tentatively accepted as new image regions, but must still pass through a *noise-cleaning* operation to eliminate the boundary and texture fragments commonly formed during splitting. This noise cleaning is a form of region growing by recursive merging.

Splitting and merging operations, which together approximate a competitive relaxation process, must be carefully balanced so that split regions are seldom immediately restored by a succession of pairwise merges. Each splitting operation provides a context for more sophisticated recursive-merging analysis; each merge, in turn, can combine related fragments that may split differently in the next round of splitting and merging.

Similar F tests are used in the two cases, but the critical levels must be different. Splitting should reach one level of detail finer than the one desired, with merging then reducing the number of subregions by about half. Gross oversplitting followed by excessive merging would take much longer and result in little if any improvement. Conservative splitting and merging, avoiding action in doubtful cases, would fail to combine the advantages of the two approaches—the overall result being undersegmentation. Slight oversplitting and remerging works best.

This competitive/cooperative balancing is very similar to that of variable selection in stepwise multiple regression. In the latter, terms are repeatedly added to or deleted from the model in a search for an optimum subset. Loops are prevented by insisting that each change increase a global criterion. There is no guarantee of reaching a global optimum unless every possible partition is tried, but recursive inclusion and deletion together offer enough freedom to achieve very good solutions.

Split/merge segmentation is more difficult because regions are two-dimensional and can combine and resplit in complex ways. Global functions of the pixel partitionings can be defined [Leclerc 88], but optimization by reassigning individual pixels is very slow on sequential machines. Local optima or flat regions in the search space are hard to escape.

The KNIFE system achieves speed and approximate optimality by taking big steps, splitting or merging large groups of pixels that appear to belong together. A statistical hypothesis test is used to determine whether each new subregion could be part of the same linear surface as one of its larger neighbors. The test permits small regions to merge easily, while larger pairs of regions may be kept separate unless the fit is quite good. A

region formed by merging two others is itself scheduled for merging.

For efficiency's sake, the KNIFE system usually accepts merges of very small subregions (e.g., four pixels) with their most statistically similar neighbors without computing linear fits and F ratios. The size threshold varies with the segmentation level and may be set or disabled by the user.

In some cases, only a single subregion may remain after all others have merged into neighboring regions. The program reassigns the original region number to this subregion and queues it for further splitting. In other cases, all of the subregions merge into neighboring regions, leaving no trace of the original region. KNIFE updates the status of any parent regions, with special attention to parents with no remaining children.

It is also possible for a region being split to be completely reformed by pairwise merging, in which case it is treated as homogeneous and is not requeued for further splitting at the current segmentation level. There might be some small advantage in backtracking and trying other thresholds or cluster groupings, but this is not now being done.

Splitting produces any number of subregions at one time, whereas merging considers only a single pair. This can lead to loops in which several small regions "swap around," repeatedly splitting from and merging with different large regions until the original configuration is repeated. Careful balancing of the critical F ratios has minimized this problem, but instances do occur. I have therefore included a loop detector in the main control sequence that declares a region homogeneous if it has been reformed from its original elements. This is rather like the human visual system noticing that an Escher print is ambiguous, rather than continuing to cycle through locally consistent interpretations.

I am not certain whether the critical F ratios for splitting and merging are unique to this program and its image domain or are more general in nature. After months of empirical testing, I determined that the best F thresholds for monochrome image segmentation have an approximate log-linear relationship across many orders of magnitude. There is also a straightforward relationship between the desired level of detail and the two F thresholds. The user need only set the level of detail; the program will select appropriate F values and adjust its search algorithm accordingly.

The statistical thresholds I use are

Seglevel:	1	2	3	4	5	6	7	8
SplitF:	12000	5800	2500	1050	420	140	45	6
MergeF:	1350	440	130	35	9	1.8	.33	.03

where Seglevel 1 represents a coarse segmentation, 3 the default level, and 5 typically a full segmentation. (Level 8 is so detailed that it should be applied only in user-specified image regions.) KNIFE often segments multiband imagery more finely than monochrome imagery for a given seglevel, so the user may want to compensate a lower setting. For precise control, the user (or control process) can single-step through a segmentation with different seglevel settings in different image areas.

KNIFE has two other parameters that vary with the segmentation level unless explicitly set by the user. *Maxnewrgns* is the number of new regions that can be formed by a thresholding operation without KNIFE's rejecting the split; this is designed to avoid wasting time on noisy splits of textured regions, but defaults to a value above 1000 subregions. *Minrgnsize* is the minimum size of regions retained by merging operations; it defaults to just a few pixels, but can be set higher to retain only large regions. (The user could also invoke a separate merging or growing step to achieve such a result.)

Most recursive segmenters maintain a tree showing the rather arbitrary pattern in which subregions were discovered. This tree is not particularly useful, however, and may have to be abandoned if a final merge pass is performed. KNIFE maintains a directed acyclic region graph, with subregions possibly belonging to more than one composite parent region. Parent/child links, which represent statistical similarity or semantic relationships, may be formed by automatic material labeling or by direct user specification. The details of updating this graph during merging are much too hairy to present in full.

For some purposes, it may be desirable to maintain a strict tree of nested regions formed at different levels of segmentation. KNIFE permits the user to specify that subregions can remerge only with one another rather than with neighboring regions. This restriction, enforced by traditional hierarchical segmenters, greatly simplifies parse representation and updating, but restricts the quality of parse found at each level of detail. Very similar effects are found in cluster analysis, where hierarchical splitting or aggregation has less freedom than an unconstrained search for the best partitioning having a particular significance level or number of clusters.

Recursive thresholding or cluster assignment, together with spatial validation, may produce subregions that are constant, textured, linearly sloped, or slowly varying. KNIFE's emphasis on threshold splitting tends toward constant regions, at least initially, but can lead to any of these surface types. The F test used in spatial validation and noise merging is specifically designed for linear surface fits. Surfaces form and reform during segmentation, with pieces detaching from one part of the region graph and attaching to another.

Because KNIFE is a reasonably fast segmenter, it could be used as a preprocessing step for either another syntactic segmenter (e.g., region clustering [Abele 80; Lumia 81, 83], graph analysis [Tanimoto 76, 77; Keng 77], or "thin plate" modeling [Leclerc 88]) or for a regionally based semantic-analysis system [Duda 70; Yakimovsky 73ab, 74; Feldman 74; Bajcsy 75; Price 76, 81; Sakai 76; Tenenbaum 76, 77; Nagao 77, 79, 80; Faugeras 79, 81, 82; Shaheen 79; Fukada 80b; Ohta 80; Levine 81ab; Wesley 82a, 86; McKeown 85; Nazif 84; Reynolds 84; Fua 86, 87ab; Bhanu 87a].

3 KNIFE Region-Growing Techniques

I have just presented a recursive splitting technique that intersperses region merging to correct errors and improve the context in which later splitting decisions are made. This section describes the reverse, a region-growing technique that intersperses region splitting to improve decision context.

Region merging typically starts with very detailed partitioning to ensure that no region will contain more than one scene object or material type. The goal is to combine these small regions into the largest possible groupings that maintain semantic purity. Growth stops when the expanding seeds collide with one another.

Region growing is more difficult because the neighbors of each growth seed need not be homogeneous. The user might trace a small area in the middle of a lake, for instance, and ask that the region be grown⁴ to its natural boundaries. We know the statistical properties of the seed region, but not those of the background populations at which growth should stop.

⁴A pixel classification approach using a priori or user-specified training signatures might be a better choice in this situation [Weymouth 83; Laws 88ab].

My KNIFE region-growing algorithm enriches the decision context by splitting neighboring regions into homogeneous subregions. Each subregion is then treated in the same manner, recursively. (The previously described splitting algorithm is used, except that noise regions can be merged only with their siblings from the same split.) If a neighboring region or subregion is small or unsplittable, the seed region tries to absorb it. Growth stops when all such neighbors are both unsplittable and unmergeable. A final step undoes the splitting, except where growth has left disconnected subregions.

Growing a single seed region is relatively simple. Each small or homogeneous neighbor is absorbed if the error of a linear surface fit does not exceed a critical F ratio. The process then begins again with the newly merged seed region and its newly computed list of neighbors. (The set of available *source regions* can be specified by either the user or the control process.)

A more difficult case arises when several seed regions share a single neighbor. KNIFE selects the seed that is most similar to the contested region, ignoring the others. If the source region is small or homogeneous, KNIFE tries to absorb it; otherwise KNIFE splits it and queues each subregion for later merging or splitting.

The most difficult case arises when many seed regions share many neighbors. This requires a global best-first merge scheduler for optimum performance, or perhaps even a search through more than just pairwise groupings. KNIFE uses a best-first merging algorithm during cluster analysis. Its integrated split/merge partitioning would be a relaxation-style solution if it were biased toward merging instead of splitting. For this application, however, KNIFE simply picks one of the source regions at random and either merges it with its most similar adjacent seed or tries to split it.

Determination of the most similar seed region is done by a histogram comparison. This is surprisingly tricky. Slow variation across regions in typical imagery makes goodness-of-fit measures impractical for comparing pixel distributions near region borders with those from region interiors. Nonparametric scores, such as the chi-square and Smirnov statistics, are especially poor because they fail to differentiate populations that are very different from those that are part of a single surface partitioned by a thresholding operation. Parametric tests, on the other hand, require unimodal and perhaps Gaussian distributions—conditions often violated by semantically meaningful scene regions and by regions created during recursive segmentation.

My heuristic solution is to sum the overlap of two smoothed histograms, then subtract a factor proportional to the difference in means of the two distributions. The resultant score is minus one for maximally separated distributions, zero for one distribution falling between the bimodal peaks of another, and one for fully overlapped distributions. The computation does not discriminate strongly against a small, tightly clustered histogram matched to a broad histogram from a much larger region.

It is easy to propose situations in which histogram-based similarity tests will fail to give the best answer. Gradient-based tests [Fennema 69; Brice 70] are equally problematical, since region borders need not lie along paths of high gradient. (Recursive segmenters have a tendency to misplace region borders by a few pixels—one of the reasons for invoking a region grower.) Local surface fits are a better approach, but depend on the size of the local operator and the geometry of the two regions. KNIFE uses region histograms to test for likely membership, full-region surface fits to verify a merge, and splitting to handle nonhomogeneous cases.

All of these techniques work in multiband imagery—whether multispectral, texture, or KNIFE's prototype-similarity bands [Laws 85]. The repeated multiband splitting is slow on current hardware, which is disconcerting because region growing is traditionally a simple, fast technique. I recommend restricting region growing to a single band whenever possible, with multiband growing or splitting used to correct errors in the final partitioning.

KNIFE's region-growing algorithm permits seed regions to bite off chunks of neighboring regions when they can't swallow their neighbors whole. This is a new technique, less well developed than KNIFE's splitting approach. Local substructure is made explicitly available during merging decisions, but KNIFE still makes only minimal use of this additional knowledge. Future versions may use classification techniques to help with the splitting and merging decisions. (Several such techniques are implemented in the KNIFE package, but are not integrated with the splitting and growing commands.) Feature extraction and object identification are so difficult that domain-specific and task-specific control processes will be needed to fully exploit these tools.

4 Future Research

There are several directions in which this work could be extended. One is to explore its relationship to relaxation segmentation algorithms [Leclerc 88] on parallel hardware, or to neural-network techniques for feature extraction and pattern recognition [Uhr 72, 82; Ballard 81; Feldman 82; Hrechanyk 82, 83; Sabbah 82; Carpenter 88; Fukushima 88]. Possibly a parallel approach using a single processor per region would be better than one with a processor per pixel, although fast histogramming and surface fitting are required. Pyramid approaches, operating simultaneously at several resolutions, might be especially effective [Uhr 72, 82; Hanson 75, 80; Arbib 76; Burt 81; Pietikäinen 81, 82b; Ahuja 84; Reynolds 84], and would eliminate some of the search loops in the current implementation.

Another direction is to develop intelligent control functions that know what to look for and how one region relates to another. This has long been a goal of region-based analysis by the image-understanding community [Guzman 67; Yakimovsky 74, 76; Wechsler 75, 77; Freuder 76a, 77; Garvey 76; Taylor 76; Ballard 77, 78; Kanade 77; Nagao 77, 79, 80; Rosenthal 78, 84; Rubin 78, 80; Russell 79; Kestner 80, 82; Selfridge 81, 82; Weymouth 81, 83; Tanimoto 82; Wesley 82b, 86; McKeown 84; Nazif 84; Reynolds 84; Harwood 87, Kohl 87]. In this regard, KNIFE may provide sufficiently powerful tools for useful region-based semantic analysis. Now that we can find regions at any level of detail, what are we to do with them?

A third direction, and the one that interests me most, is the integration of additional low-level techniques to improve performance. Given a generic task, a partially interpreted image, and a set of techniques, how can we find the best partitioning or object labeling in the least amount of time?

Representing the generic task is the most difficult part. I envision expert-system rule bases or procedural control algorithms for cloud cover estimation, road following, target cuing, angiogram analysis, and the like, but I do not expect the full analysis to be accomplished at the level of the image segmenter. KNIFE should interface with high-level systems that do evidential reasoning and dynamic replanning, but it need not have such capabilities itself.

The KNIFE toolbox already includes segmentation, growing, merging, cluster analysis,

and classification. There are many ways that these can be combined to form more powerful operators. For instance, classification is very good at finding additional objects once a prototype has been identified; I have implemented a temporal tracking system based on this principle. (More intelligence is needed, however, in updating spectral signatures as tracking progresses, as well as in handling label inheritance during region merging.)

Edge detection and other gradient-based techniques should be added next [Jarvis 75; Irwin 84; Nazif 84; Reynolds 84; Belknap 86]. When splitting a region, for instance, it makes sense to concentrate initial effort along ridges of high gradient. Statistics gathered over the region could be used to customize an edge detector for optimum performance. Statistics from two regions can help to position the boundary between them precisely. (I have investigated the use of multivariate maximum-likelihood tests based on smoothed multinomial distributions for this purpose, but I must incorporate local surface slope as well.)

KNIFE should also be extended to handle higher-order surface fits. The current program cannot represent geometric primitives such as cylinders and spheres, and so cannot adequately parse images from the industrial-inspection domain. Rounded objects, such as bushes and trees in range imagery, are also problematic. KNIFE should incorporate RANSAC-style techniques [Bolles 81; Fischler 81b] for recognizing, describing, and delineating these curved surfaces, making appropriate use of range data or other bands with special characteristics.

Flat surfaces seen in perspective offer a similar challenge. Textures such as grass and gravel change in deterministic ways from the foreground to the horizon [Witkin 80, 81ab; Pentland 84, 86ab; Strat 86]. We have begun to study these effects, including ways to estimate camera and illumination models from single-view image data, but we have not integrated these techniques with image segmentation. Even simple texture segmentation, without perspective distortions, remains a problem, since current texture operators give misleading responses near object borders and for very small objects.

Even the capabilities currently in KNIFE could be improved. The segmenter has difficulty with diagonal edges, breaking out the mixed pixels as separate regions and then remerging them. Switching to an eight-connected or perhaps hexagonal neighborhood definition would be one way to combat this inefficiency.

Likewise, the median splitting described above could be replaced by a fast region-growing technique [Somerville 76; Yakimovsky 76], possibly using low-gradient image areas as seed regions. Classification could be pursued more vigorously as a way of locating discontiguous objects or material classes, as in land-use classification [Swain 68, 81; Sadjadi 79; Connors 82, 84; Cate 83; Trivedi 84c, 85; Harlow 85], target cuing [Aggarwal 78abc; Lutton 80; Trivedi 84ab], blob extraction [Panda 77, 78; Blanz 81; CHChen 81; Skevington 81; Hartley 82], cell or particle counting, digital character recognition, etc. Other techniques in the current package could be streamlined by adding more sophisticated control strategies that reason about goals and partitioning failures. Such optimizations within the modular KNIFE package are not too difficult to implement.

5 Examples

KNIFE is a general-purpose segmenter, and diverse examples are required to illustrate its capabilities. The program's data structures have not been extended to curved or

polynomial surfaces, so I shall focus on domains where linear surfaces are typical. I intend to show KNIFE's strengths and weaknesses under realistic operating conditions.

The images shown below range from 64×64 to 256×256 pixels, although the current C implementation of KNIFE can be used on images as large as 512×512 .⁵ Segmentation of 256×256 images takes from five minutes to several hours, depending on scene content and the requested level of detail: typically 15 minutes for segmentation level 1, 30 minutes for level 3, and 60 minutes for level 5 on a VAX 11/780. Multiband images take proportionately longer unless the scene regions are easily segmented in each band.

Figure 1(a) shows a house and yard with circular and side driveways, a planter and pool area, grass, trees, and bushes. There are small but quite definite shadows from the house and chimneys, trees, and one piece of pool furniture. The roof has some texture, but not enough to justify use of even a 3×3 texture operator.

Figures 1(b) through (d) depict three levels of segmentation. Seglevel 1 is insufficient to split the image, which is not too surprising if this 80×87 image is viewed as a detail in some larger scene. (Seglevel 1 is intended for crude partitioning, such as separating sky from ground. KNIFE does not model the human tendency to find regions in proportion to image size.) Seglevel 3 works well, although it misses the house shadow and the deep end of the pool. Seglevel 5 finds much of the roof structure, but is not quite strong enough to extract trees in the front lawn. (The trees are perceived primarily through shadow cues that require sophisticated surface modeling.) Additional semantic features could be detected at higher seglevel values, although noise regions would also be created.

Figure 2 shows a similar sequence for a set of apartment buildings. Seglevel 1 pulls out only the shadows, as if these were dark buildings next to white parking lots. Seglevel 3 does a fair job of extracting buildings, streets, and sidewalks, as well as a partially obscured paved court among the trees. Seglevel 5 finds cars and other small structures, but semantic filtering is needed to suppress the corresponding arboreal detail. (KNIFE can use texture only to find additional detail rather than to discount visible structure in the intensity band. A higher-level control process could prevent fragmentation of homogeneously textured regions, as was done in the Ohlander segmenter, but only at the risk of missing important details hidden within the trees.)

Figure 3(a) is a well-known image of Fort Belvoir. The tree texture is clearly visible at this resolution, but is challenging for even a 3×3 texture operator, such as the log variance in Figure 3(b). I was tempted to segment a higher-resolution version of this image—a common dodge in texture research—but wanted to show what KNIFE could do with textures that humans commonly exploit.

KNIFE is able to partition the two-band image at Seglevel 3, although it takes several hours because of thresholding fragmentation. (I have to increase the default maxnewrgns parameter, which controls KNIFE's mechanism for limiting CPU time wasted on splitting textured regions.) The results in Figure 3(c) are unimpressive, but better than the complete failure experienced when no texture band is used. Note, however, that the intensity band alone is able to produce the Seglevel 5 partitioning in Figure 3(d). I often find that increasing seglevel has much the same effect as adding data bands.

Figure 4 illustrates a domain in which sampling rate has exceeded optical resolution; and texture measures are likely to tell us more about the sensor than about the scene. The image is reportedly an early FLIR (forward-looking infrared) picture of an armored tank.

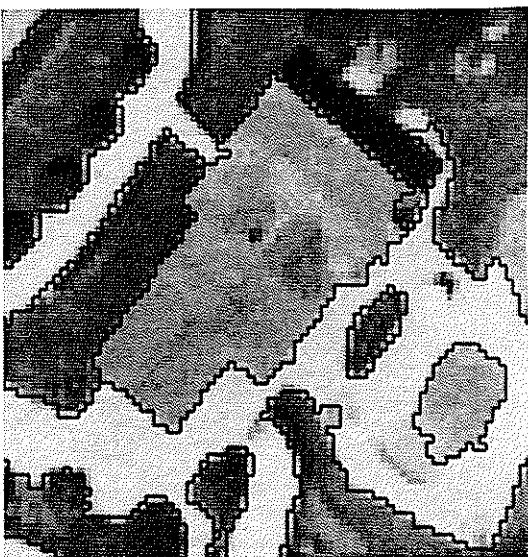
⁵Larger images might cause register overflow in its cluster-analysis and surface-fit routines.



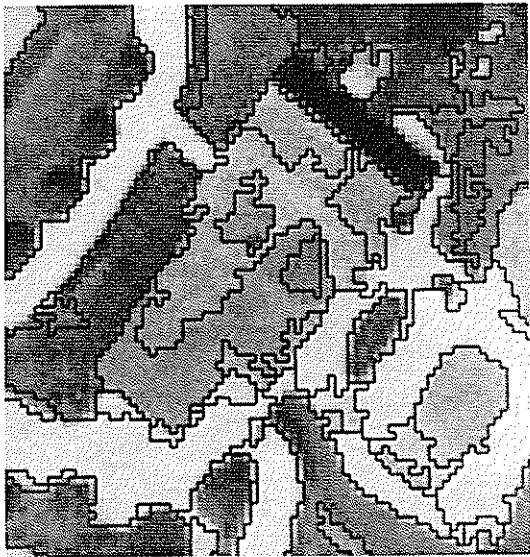
(a) Intensity band.



(b) Level 1 (1 region).



(c) Level 3 (18 regions).



(d) Level 5 (77 regions).

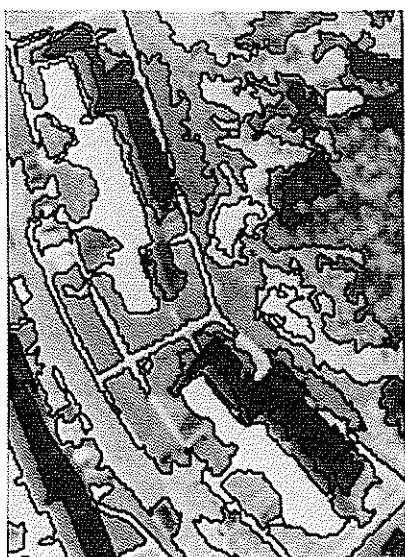
Figure 1: Segmentation of monochrome HOUSE image.



(a) Intensity Band



(b) Level 1 (7 Regions)



(c) Level 3 (62 Regions)



(d) Level 5 (607 Regions)

Figure 2: Segmentation of Monochrome BUILDING Image



(a) Intensity (I) Band



(b) Log-Variance ($LV3$) Band

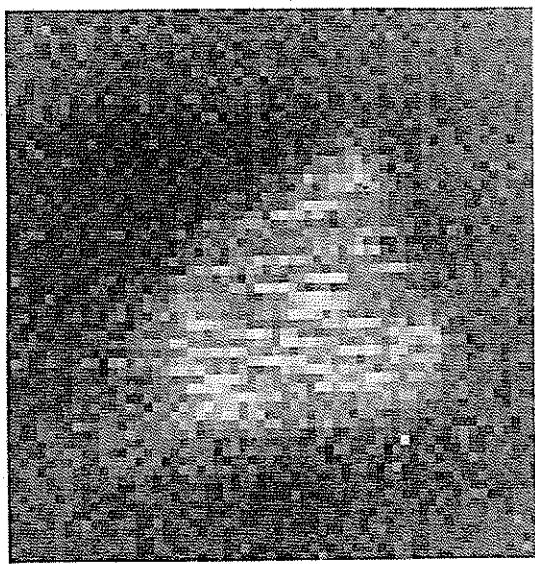


(c) Level 3, $I+LV3$ (104 Regions)

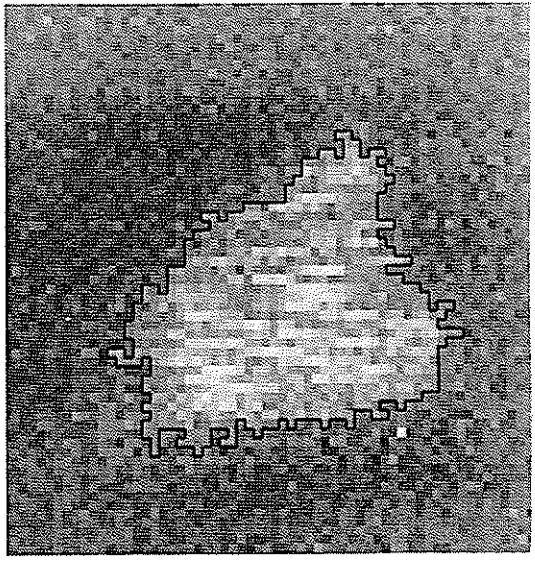


(d) Level 5, I only (1073 Regions)

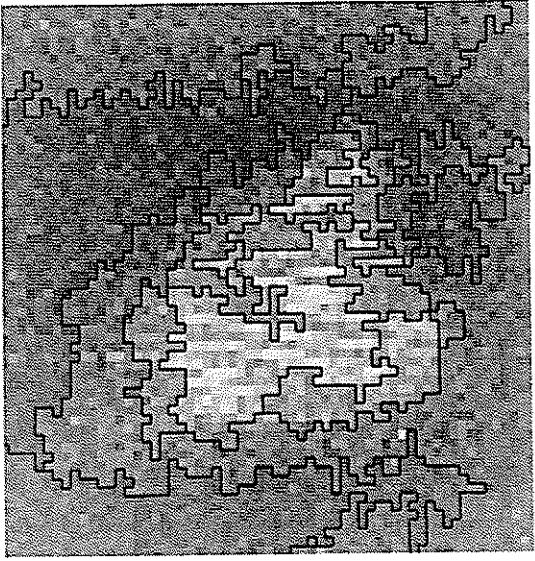
Figure 3: Segmentation of Textured FORT BELVOIR Image



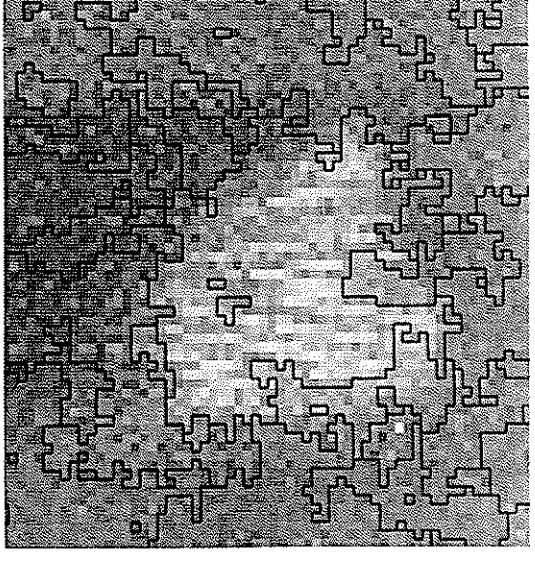
(a) Infrared Intensity (I) Band



(b) Level 3, $I+LV3$ (2 Regions)



(c) Level 5, I Only (5 Regions)



(d) Level 5, $I+LV3$ (78 Regions)

Figure 4: Segmentation of Textured TANK Image

KNIFE is unable to locate the hotspot at Seglevel 1. At Seglevel 3, it extracts the blob in Figure 4(b) from either intensity alone or intensity plus local log variance. Figures 4(c) and (d) show the Seglevel 5 results when intensity alone is used or the texture measure is made available. Neither of these partitionings closely matches my own perception of the image substructure.

Figure 5 shows a simple road scene taken from a prototype of the FMC autonomous vehicle.⁶ Figure 5(a) is actually the vividness band of my VHS transform, but differs little from other intensity or perceptual brightness bands. The Seglevel 1 result would be useful for road following, although bright regions of the mountain face are merged with the sky. Seglevel 3 corrects this problem, but begins to extract substructure in the road. (A higher-level control process should be able to sew the pieces back together if boundary smoothness and perhaps color characteristics are now computed for each region.) Seglevel 5 is useful only if one is looking for potholes or distant objects.

Figures 6(b) and (c) are the hue and saturation bands for this image. Hue is of little use beyond distinguishing sky from ground—although that does clear up our problem with the mountain faces. Saturation appears to have more useful structure, but is correlated with the vividness band. Level 1 and 3 segmentations are slightly more detailed than their intensity-only counterparts and have better spatial coherence.

Figures 7 and 8 illustrate this same analysis for a more difficult image.⁷ Monochrome segmentations are quite good, especially for level 3, although the foreground telephone pole is not extracted cleanly and the sky shows sensor-induced banding. Color partitioning is even better, although it still doesn't get the bottom half of the telephone pole. (Note that the bright top half in the hue image is an ACHROMATIC region. Bright sky altered sensor response to the dark telephone pole.) The oddly shaped shadow regions may be difficult to interpret, but a task-independent segmenter can hardly be expected to do better.

Figures 9 and 10 illustrate KNIFE's region-growing capability. The first image is monochrome, although identical results are obtained with VHS input. I took the bottom portion of the image as a seed and grew it with the seglevel set to 1. Growth was stopped by the pole of a "No Parking" sign, or perhaps because a neighboring region could not be split at Seglevel 1. Monochrome growth worked well because the road is homogeneous and very different from its surroundings, but would have failed if the road region touched the sky. (This "leakage" or overmerging did occur in another image from the same sequence.) Undermerging is also possible if monochrome segmentation is not powerful enough to split neighboring regions.

The second image, Figure 10, shows growth of a color region. KNIFE tends to see vertical banding in the road, as shown in Figure 5(c). The seed region and growing algorithm are sufficient to find the left and center road portions, but the surface fit test then rejects most of the right side for inclusion in the same region. The full road is extracted if only the vividness band is used, but hue and saturation combine with the slight intensity differences to block the linear-surface fit. Extension to curved or polynomial surface models would solve this problem.

⁶The image name, ALV519, refers to a frame-sequence number.

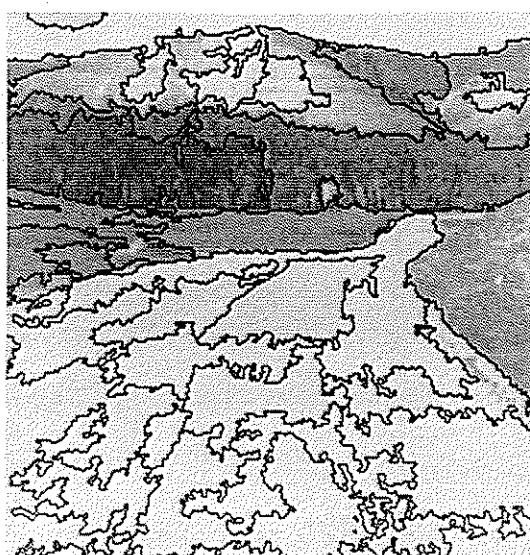
⁷It is especially difficult for edge-based segmentation methods because of road texture, tree shadows, and interlace jitter. Region-growing and classification approaches also have trouble with the tree shadows.



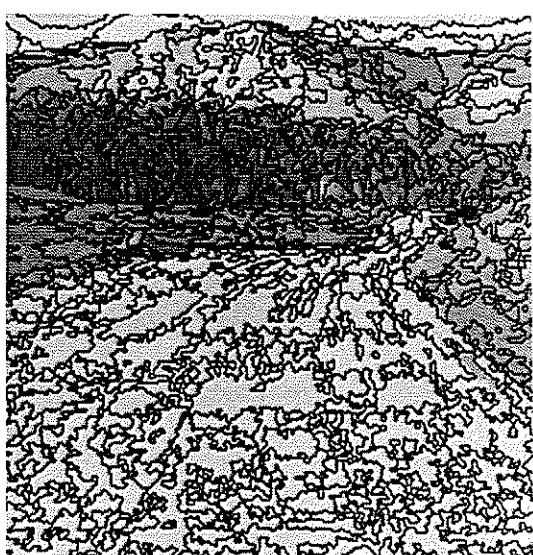
(a) Intensity (V) Band



(b) Level 1 (10 Regions)



(c) Level 3 (52 Regions)



(d) Level 5 (834 Regions)

Figure 5: Segmentation of Monochrome ALV519 Image

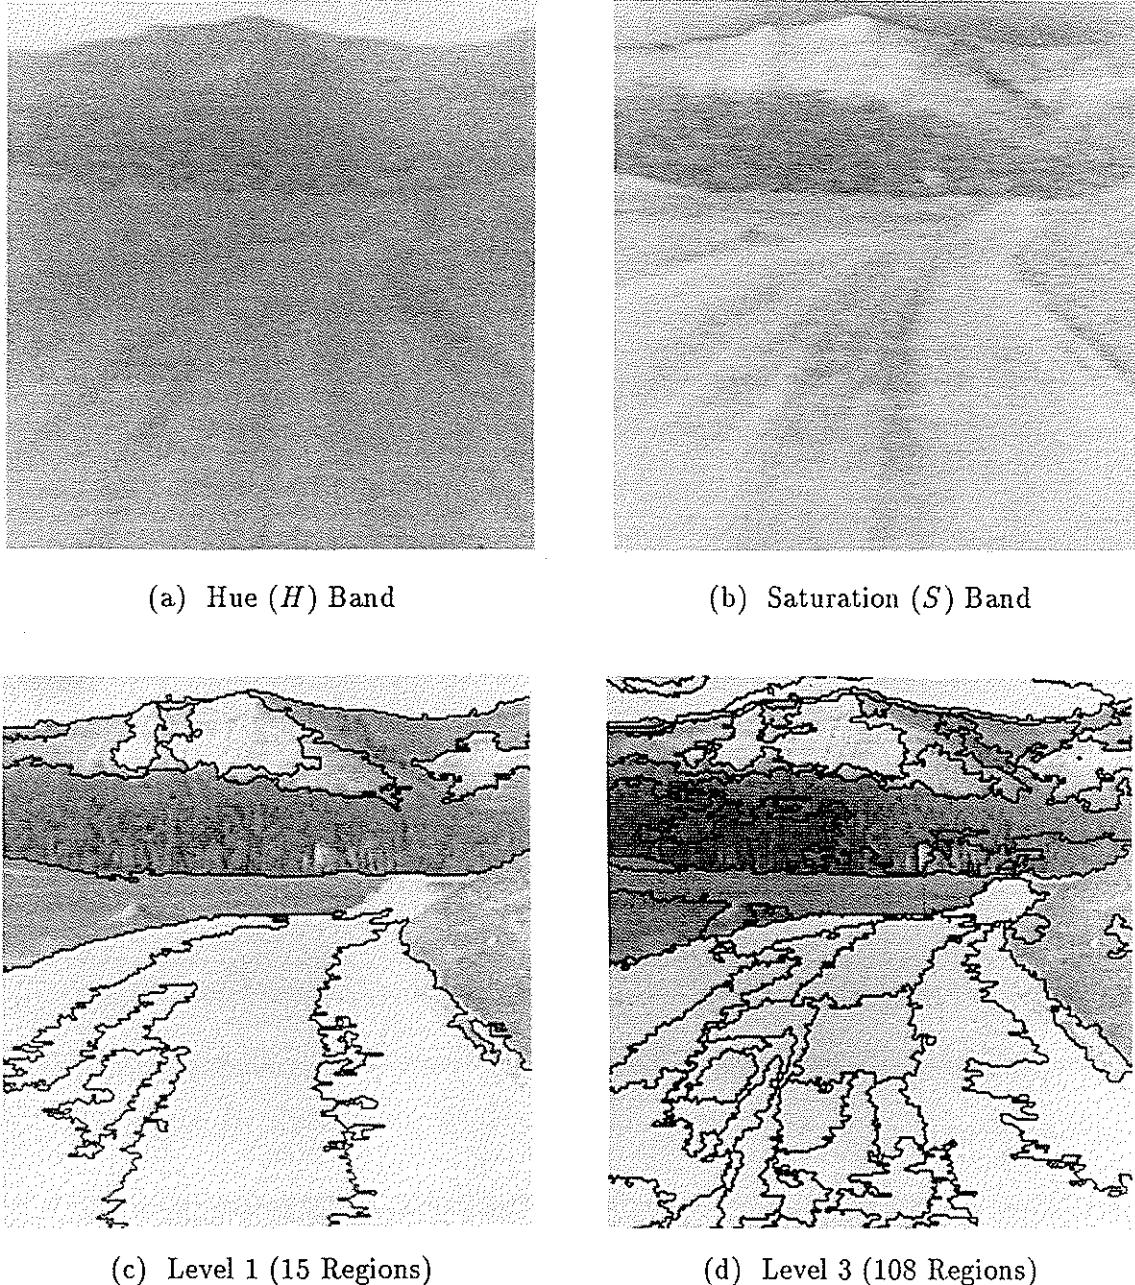
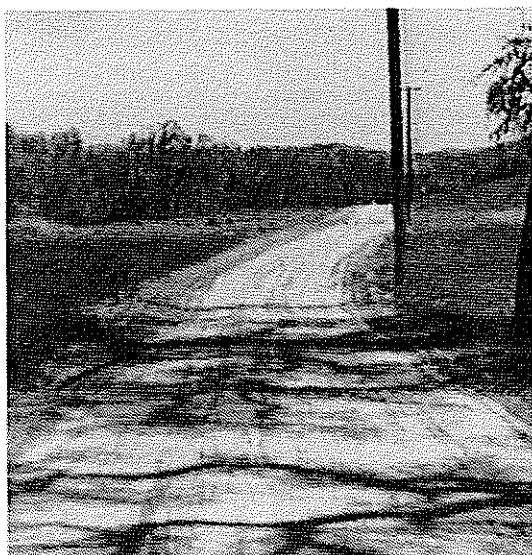
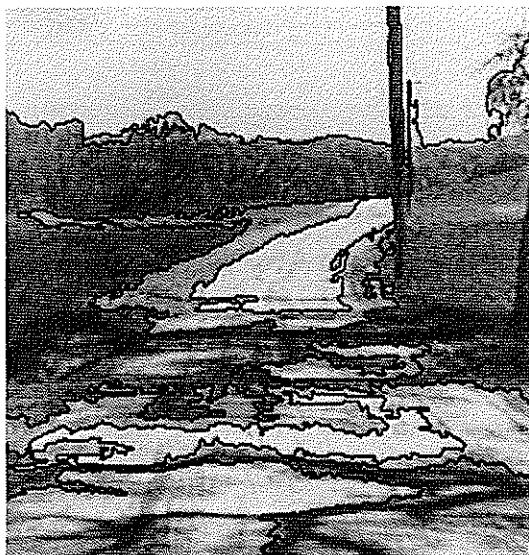


Figure 6: Segmentation of Color ALV519 Image



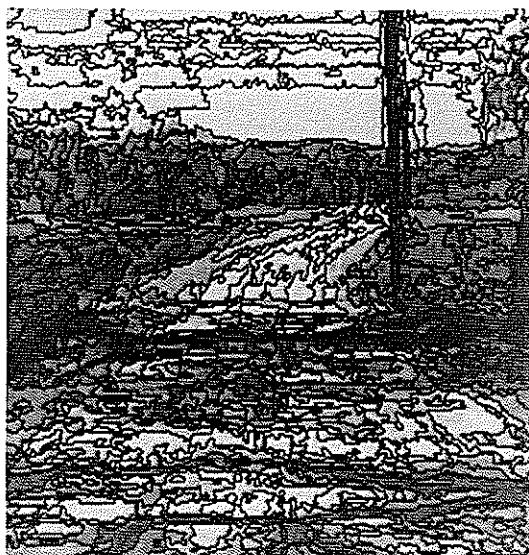
(a) Intensity (V) Band



(b) Level 1 (10 Regions)



(c) Level 3 (85 Regions)



(d) Level 5 (952 Regions)

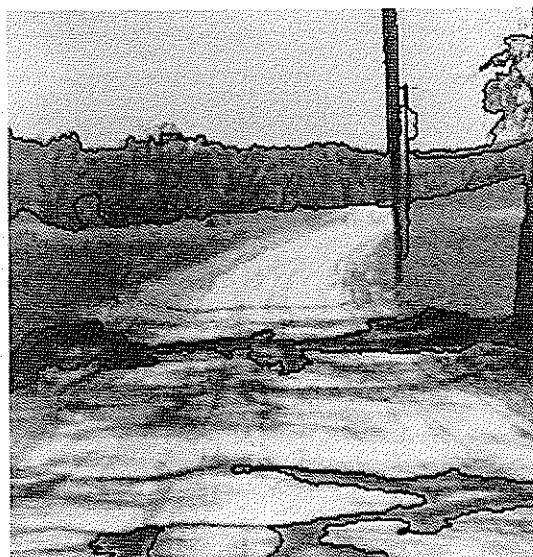
Figure 7: Segmentation of Monochrome ALV533 Image



(a) Hue (H) Band



(b) Saturation (S) Band

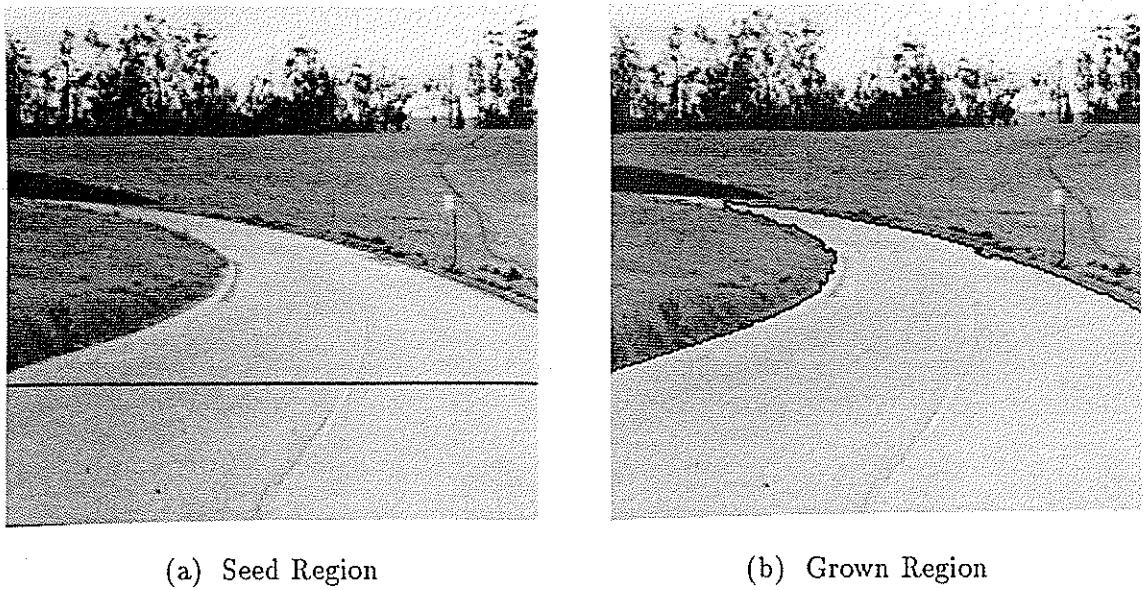


(c) Level 1 (10 Regions)



(d) Level 3 (109 Regions)

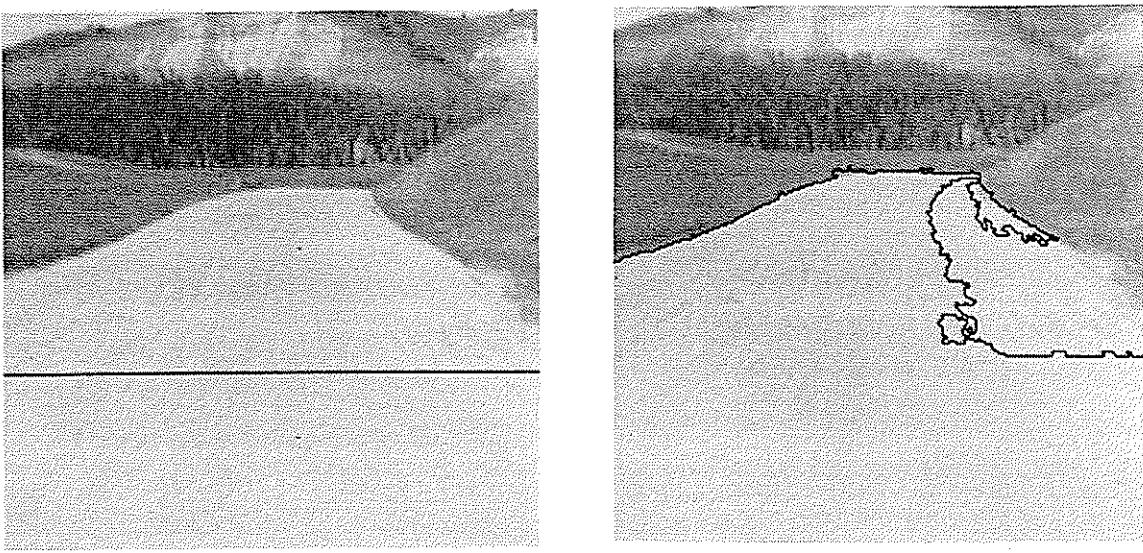
Figure 8: Segmentation of Color ALV533 Image



(a) Seed Region

(b) Grown Region

Figure 9: Region Growing in Monochrome ROAD Image



(a) Seed Region (V Band)

(b) Grown Region

Figure 10: Region Growing in Color ALV500 Image

6 Summary

I have described two capabilities of the KNIFE feature-extraction package: integrated split/merge segmentation and split/merge region growing. These exploit new techniques for color, texture, and prototype-similarity transformation; histogram analysis; multiband cluster analysis; connected-component extraction; statistical goodness-of-fit comparisons; and sloped-region merging for noise cleaning and spatial verification—all tied together by a sophisticated command language, bookkeeping system, and object-oriented programming environment. Although many improvements are possible, the current system is relatively fast, easy to use, and powerful. Its segmentations of monochrome imagery rival those of Ohlander-style segmenters using color imagery; with multiband input, KNIFE can do even better.

Acknowledgment

This work was supported by Defense Advanced Research Projects Agency under Contract Nos. MDA903-86-C-0084 and DACA76-85-C-0004.

References

[Abele 80] L.B. Abele, "Feature Selection by Space Invariant Comparisons with Applications to the Segmentation of Textured Pictures," *Proc. 5th Int. Jnt. Conf. on Pattern Recognition*, Miami Beach, Florida, pp. 535-539 (1-4 December 1980).

[Agin 79] G.J. Agin, *Hierarchical Representation of Three-Dimensional Objects Using Verbal Models*, Technical Note 182, Artificial Intelligence Center, SRI International, Menlo Park, California (March 1979).

[Aggarwal 78a] R.K. Aggarwal, "Image Segmentation Using Prototype Similarity," *Proc. SPIE* (March 1978).

[Aggarwal 78b] R.K. Aggarwal, "Adaptive Image Segmentation Using Prototype Similarity," *Proc. 4th IEEE Conf. on Pattern Recognition and Image Processing*, Chicago, pp. 354-359 (May 1978).

[Aggarwal 78c] R.K. Aggarwal, "Image Segmentation for Syntactic Classification of Large Images," *Proc. DARPA Image Understanding Workshop*, Cambridge, Massachusetts, pp. 125-127 (3-4 May 1978).

[Ahuja 84] N. Ahuja and S. Swamy, "Multiprocessor Pyramid Architectures for Bottom-Up Image Analysis," *IEEE Trans. on Pattern Analysis and Machine Intelligence*, Vol. PAMI-6, No. 4, pp. 463-475 (July 1984).

[Arbib 76] M.A. Arbib and E.M. Riseman, "Computational Techniques in Visual Systems," Parts I and II, COINS Technical Reports 76-10 and 76-11, University of Massachusetts at Amherst, Amherst, Massachusetts (July 1976).

[Asano 81] T. Asano and N. Yokoya, "Image Segmentation Schema for Low-Level Computer Vision," *Pattern Recognition*, Vol. 14, Nos. 1-6, pp. 267-273 (1981).

[Bajcsy 74] R. Bajcsy and L.I. Lieberman, "Computer Description of Real Outdoor Scenes," *Proc. 2nd Int. Jnt. Conf. on Pattern Recognition*, Copenhagen, Denmark, pp. 174-179 (13-15 August 1974).

[Bajcsy 75] R. Bajcsy and M. Tavakoli, "Image Filtering—A Context Dependence Process," *IEEE Trans. on Circuits and Systems*, pp. 463-474 (May 1975).

[Ball 65] G.H. Ball and D.J. Hall, *ISODATA, A Novel Method of Data Analysis and Pattern Classification*, Stanford Research Institute, Menlo Park, California (April 1965).

[Ball 67] G.H. Ball and D.J. Hall, "A Clustering Technique for Summarizing Multivariate Data," *Behavioral Science*, Vol. 12, pp. 153-155 (March 1967).

[Ballard 77] D.H. Ballard, C.M. Brown, and J.A. Feldman, "An Approach to Knowledge-Directed Image Analysis," *Proc. 5th Int. Jnt. Conf. on Artificial Intelligence*, Cambridge, Massachusetts, pp. 664-670 (22-25 August 1977). Reprinted in [Pao 82], pp. 546-552.

[Ballard 78] D.H. Ballard, "Model-Directed Detection of Ribs in Chest Radiographs," *Proc. 4th Int. Jnt. Conf. on Pattern Recognition*, Kyoto, Japan, pp. 907-910 (November 1978). Also published as TR11, Computer Science Department, University of Rochester (March 1978).

[Ballard 81] D.H. Ballard, "Parameter Networks: Towards a Theory of Low Level Vision," *Proc. 7th Int. Jnt. Conf. on Artificial Intelligence*, Vancouver, British Columbia, Canada, pp. 1068-1078 (24-28 August 1981).

[Barrett 81] W.A. Barrett, "An Iterative Algorithm for Multiple Threshold Detection," *Proc. IEEE Conf. on Pattern Recognition and Image Processing*, Dallas, Texas, pp. 273-278 (August 1981).

[Beaulieu 83] J.-M. Beaulieu and M. Goldberg, "Step-wise Optimization for Hierarchical Picture Segmentation," *Proc. IEEE Conf. on Computer Vision and Pattern Recognition*, Washington, D.C., pp. 59-64 (19-23 June 1983).

[Belknap 86] R. Belknap, E. Riseman, and A. Hanson, "The Information Fusion Problem and Rule-Based Hypotheses Applied to Complex Aggregations of Image Events," *Proc. IEEE Conf. on Computer Vision and Pattern Recognition*, Miami Beach, Florida, pp. 227-234 (22-26 June 1986). Also published in *Proc. DARPA Image Understanding Workshop*, Miami Beach, Florida, pp. 279-292 (9-10 December 1985).

[Bhanu 82] B. Bhanu and O.D. Faugeras, "Segmentation of Images Having Unimodal Distributions," *IEEE Trans. on Pattern Analysis and Machine Intelligence*, Vol. PAMI-4, No. 4, pp. 408-419 (July 1982).

[Bhanu 87a] B. Bhanu and P. Symosek, "Interpretation of Terrain Using Hierarchical Symbolic Grouping for Multi-Spectral Images," *Proc. DARPA Image Understanding Workshop*, Los Angeles, California, pp. 466-474 (23-25 February 1987).

[Bhanu 87b] B. Bhanu and B.A. Parvin, "Segmentation of Natural Scenes," *Pattern Recognition*, Vol. 20, No. 5, pp. 487-496 (1987).

[Binford 82] T.O. Binford, "Survey of Model Based Image Analysis Systems," *Int. J. of Robotics Research*, Vol. 1, No. 1, pp. 18-64, (Spring 1982).

[Blanz 81] W.E. Blanz and E.R. Reinhardt, "Image Segmentation by Pixel Classification," *Pattern Recognition*, Vol. 13, No. 4, pp. 293-298 (1981).

[Bolles 78] R.C. Bolles, L.H. Quam, M.A. Fischler, and H.C. Wolf, "The SRI Road Expert: Image-to-Database Correspondence," *Proc. DARPA Image Understanding Workshop*, Pittsburgh, Pennsylvania, pp. 163-174 (14-15 November 1978).

[Bolles 81] R.C. Bolles, "A RANSAC-Based Approach to Model Fitting and its Application to Finding Cylinders in Range Data," *Proc. 7th Int. Conf. on Artificial Intelligence*, Vancouver, British Columbia, Canada, pp. 637-643 (24-28 August 1981). Also published as Technical Note 235, Artificial Intelligence Center, SRI International, Menlo Park, California (March 1981).

[Bolles 82] R.C. Bolles and R.A. Cain, "Recognizing and Locating Partially Visible Objects: The Local-Feature-Focus Method," *Int. J. of Robotics Research*, Vol. 1, No. 3, pp. 57-82 (Fall 1982).

[Bolles 86] R.C. Bolles and P. Horaud, "3DPO: A Three-Dimensional Part Orientation System," *Int. J. of Robotics Research*, Vol. 5, No. 3, pp. 3-26 (1986).

[Brice 70] C.R. Brice and C.L. Fennema, "Scene Analysis Using Regions," *Artificial Intelligence*, Vol. 1, No. 3, pp. 205-226 (Fall 1970). Also published as Technical Note 17, Artificial Intelligence Group, Stanford Research Institute, Menlo Park, California (April 1970).

[Brooks 79] R.A. Brooks, R. Greiner, and T.O. Binford, "The ACRONYM Model-Based Vision System," *Proc. 6th Int. Conf. on Artificial Intelligence*, Tokyo, Japan pp. 105-113 (20-23 August 1979).

[Brooks 81] R.A. Brooks, *Symbolic Reasoning Among 3-D Models and 2-D Images*, Ph.D. Thesis, STAN-CS-81-861, Stanford University, Stanford, California (June 1981). Summarized in *Artificial Intelligence Journal*, Vol. 17, No. 1-3, pp. 385-348 (August 1981).

[Brooks 84] R.A. Brooks, *Model-Based Computer Vision*, UMI Research Press, Ann Arbor, Michigan (1984).

[Browning 82] J.D. Browning and S.L. Tanimoto, "Segmentation of Pictures into Regions with a Tile-by-Tile Method," *Pattern Recognition*, Vol. 15, No. 1, pp. 1-10 (1982).

[Bryant 79] J. Bryant, "On the Clustering of Multidimensional Pictorial Data," *Pattern Recognition*, Vol. 11, pp. 115–125 (1979).

[Burt 81] P. Burt, T.H. Hong, and A. Rosenfeld, "Segmentation and Estimation of Image Region Properties through Cooperative Hierarchical Computation," *IEEE Trans. on Systems, Man, and Cybernetics*, Vol. SMC-11, No. 12 (December 1981). Also published as TR-927, Computer Vision Laboratory, Computer Science Center, University of Maryland, College Park, Maryland (August 1980).

[Carlotto 85] M.J. Carlotto, "Histogram Analysis using a Scale-Space Approach," *Proc. IEEE Conf. on Computer Vision and Pattern Recognition*, San Francisco, pp. 334–340 (19–23 June 1985).

[Carlton 77] S.G. Carlton and O.R. Mitchell, "Image Segmentation Using Texture and Gray Level," *Proc. DARPA Image Understanding Workshop*, Minneapolis, Minnesota, pp. 71–77 (April 1977). Also in *IEEE Conf. on Pattern Recognition and Image Processing*, Troy, New York, pp. 387–391 (6–8 June 1977).

[Carpenter 88] G.A. Carpenter and S. Grossberg, "The ART of Adaptive Pattern Recognition by a Self-Organizing Neural Network," *IEEE Computer*, Vol. 21, No. 3, pp. 77–88 (March 1988).

[Cate 83] R.B. Cate, T.B. Dennis, J.T. Malin, K.S. Nedelman, M.H. Trenchard, and R.M. Bizzell, "A New Approach to Extraction of Invariant Scene Characteristics," *Proc. Trends and Applications*, Gaithersburg, Maryland, pp. 210–215 (25–26 May 1983).

[CHChen 81] C.H. Chen, "On the Statistical Image Segmentation Techniques," *Proc. IEEE Conf. on Pattern Recognition and Image Processing*, Dallas, Texas, pp. 262–266 (August 1981).

[PCChen 79] P.C. Chen and T. Pavlidis, "Segmentation by Texture Using a Co-Occurrence Matrix and a Split-and-Merge Algorithm," *Computer Graphics and Image Processing*, Vol. 10, pp. 172–182 (1979). Also published as TR-237, Princeton University, Princeton, New Jersey (January 1978).

[PCChen 80] P.C. Chen and T. Pavlidis, "Image Segmentation as an Estimation Problem," *Computer Graphics and Image Processing*, Vol. 12, pp. 153–172 (1980).

[PCChen 83] P.C. Chen and T. Pavlidis, "Segmentation by Texture Using Correlation," *IEEE Trans. on Pattern Analysis and Machine Intelligence*, Vol. PAMI-5, No. 1, pp. 64–69 (1983). Also published as *Proc. 5th Int. Jnt. Conf. on Pattern Recognition*, Miami Beach, Florida, pp. 551–553 (December 1980).

[Coleman 77] G.B. Coleman, *Image Segmentation by Clustering*, USCIPI Report 750, Image Processing Institute, University of Southern California, Los Angeles (July 1977).

[Coleman 79] G.B. Coleman and H.C. Andrews, "Image Segmentation by Clustering," *Proc. IEEE*, Vol. 67, No. 5, pp. 773–785 (May 1979).

[Conners 82] R.W. Conners, M.M. Trivedi, and C.A. Harlow, *The Segmentation of a Complex Urban Scene Using Texture Operators*, Remote Sensing and Image Processing Technical Report 401.82, Dept. of Electrical and Computer Engineering, Louisiana State University (1982).

[Conners 84] R.W. Conners, M.M. Trivedi, and C.A. Harlow, "Segmentation of a High-Resolution Urban Scene Using Texture Operators," *Computer Vision, Graphics, and Image Processing*, Vol. 25, pp. 273–310 (March 1984).

[Danker 81] A.J. Danker and A. Rosenfeld, "Blob Detection by Relaxation," *IEEE Trans. on Pattern Analysis and Machine Intelligence*, Vol. PAMI-3, No. 1, pp. 79–92 (January 1981).

[Davis 82] L.S. Davis and A. Mitiche, "MITES: A Model-Driven, Iterative Texture Segmentation Algorithm," *Computer Graphics and Image Processing*, Vol. 19, pp. 95–110 (1982).

[Deguchi 78] K. Deguchi and I. Morishita, "Texture Characterization and Texture-Based Image Partitioning Using Two-Dimensional Linear Estimation Techniques," *IEEE Trans. on Computers*, Vol. C-27, No. 8, pp. 739-745 (August 1978).

[Doherty 86] M.F. Doherty, C.M. Bjorklund, and M.T. Noga, "Split-Merge-Merge: An Enhanced Segmentation Capability," *Proc. IEEE Conf. on Computer Vision and Pattern Recognition*, Miami Beach, Florida, pp. 325-330 (22-26 June 1986).

[Duda 70] R.O. Duda, *Some Current Techniques for Scene Analysis*, Technical Note 46, Artificial Intelligence Center, Stanford Research Institute, Menlo Park, California (October 1970).

[Faugeras 79] O.D. Faugeras and M. Berthod, "Scene Labeling: An Optimization Approach," *Proc. IEEE Conf. on Pattern Recognition and Image Processing*, Chicago, Illinois, pp. 318-326 (August 1979).

[Faugeras 81] O.D. Faugeras and K. Price, "Semantic Description of Aerial Images Using Stochastic Labeling," *IEEE Trans. on Pattern Analysis and Machine Intelligence*, Vol. PAMI-3, No. 6, pp. 633-642 (November 1981). Also in *Proc. DARPA Image Understanding Workshop*, College Park, Maryland, pp. 89-94 (30 April 1980), and in *Proc. 5th Int. Jnt. Conf. on Pattern Recognition*, Miami Beach, Florida, pp. 352-357 (December 1980).

[Faugeras 82] O.D. Faugeras, "Relaxation Labeling and Evidence Gathering," *Proc. IEEE Conf. on Pattern Recognition and Image Processing*, Las Vegas, Nevada, pp. 672-677 (14-17 June 1982).

[Feldman 74] J.A. Feldman and Y. Yakimovsky, "Decision Theory and Artificial Intelligence: I. A Semantics-Based Region Analyzer," *Artificial Intelligence*, Vol. 5, pp. 349-371 (Winter 1974).

[Feldman 82] J.A. Feldman, "Four Frames Suffice: A Provisionary Model of Vision and Space," *University of Massachusetts Workshop on Vision, Brain and Cooperative Computation*, Amherst, Massachusetts (May 1983). Also published as TR99, Computer Science Department, University of Rochester (September 1982).

[Feng 75] H.Y.F. Feng and T. Pavlidis, "The Generation of Polygonal Outlines of Objects from Gray Level Pictures," *IEEE Trans. on Circuits and Systems*, Vol. CAS-22, No. 5, pp. 427-439 (May 1975). Also published as TR-150, Computer Science Laboratory, Department of Electrical Engineering, Princeton University, Princeton, New Jersey (April 1974).

[Fennema 69] C.L. Fennema and C.R. Brice, *A Region-Oriented Data Structure*, Technical Note 7, Artificial Intelligence Group, Stanford Research Institute, Menlo Park, California (May 1969).

[Fischler 81a] M.A. Fischler, J.M. Tenenbaum, and H.C. Wolf, "Detection of Roads and Linear Structures in Low-Resolution Aerial Imagery Using a Multisource Knowledge Integration Technique," *Computer Graphics and Image Processing*, Vol. 15, No. 3, pp. 201-223 (March 1981).

[Fischler 81b] M.A. Fischler and R.C. Bolles, "Random Sample Consensus: A Paradigm for Model Fitting with Applications to Image Analysis and Automated Cartography," *Communications of the ACM*, Vol. 24, No. 6, pp. 381-395 (June 1981). Also published in *Proc. DARPA Image Understanding Workshop*, College Park, Maryland, pp. 71-88 (30 April 1980), and as Technical Note 213, Artificial Intelligence Center, SRI International, Menlo Park, California (March 1980).

[Fischler 82] M.A. Fischler, S.T. Barnard, R.C. Bolles, M. Lowry, L. Quam, G. Smith, and A. Witkin, "Modeling and Using Physical Constraints in Scene Analysis," *Proc. National Conf. on Artificial Intelligence*, Pittsburgh, Pennsylvania, pp. 30-35 (18-20 August 1982). Also published as Technical Note 267, Artificial Intelligence Center, SRI International, Menlo Park, California (September 1982).

[Freuder 76a] E.C. Freuder, *A Computer System for Visual Recognition Using Active Knowledge*, Ph.D. Thesis, AI-TR-345, Massachusetts Institute of Technology, MIT AI Laboratory, Cambridge, Massachusetts (June 1976). Summarized in *Proc. 5th Int. Jnt. Conf. on Artificial Intelligence*, Cambridge, Massachusetts, pp. 671-677 (22-25 August 1977), and reprinted in [Pao 82], pp. 511-517.

[Freuder 76b] E. Freuder, "Affinity: A Relative Approach to Region Growing," *Computer Graphics and Image Processing*, Vol. 5, No. 2, pp. 254-264 (June 1976).

[Freuder 77] E.C. Freuder, "A Computer System for Visual Recognition Using Active Knowledge," *Proc. 5th Int. Jnt. Conf. on Artificial Intelligence*, Cambridge, Massachusetts, pp. 671-677 (1977). Reprinted in [Pao 82], pp. 511-517.

[Fua 86] P. Fua and A.J. Hanson, *Using Generic Geometric Knowledge to Delineate Cultural Objects in Aerial Imagery*, Technical Note 378, Artificial Intelligence Center, SRI International, Menlo Park, California (March 1986).

[Fua 87a] P. Fua and A.J. Hanson, "Resegmentation Using Generic Shape: Locating General Cultural Objects," *Pattern Recognition Letters*, Vol. 5, pp. 243-252 (1987).

[Fua 87b] P. Fua and A.J. Hanson, "Using Generic Geometric Models for Intelligent Shape Extraction," *Proc. DARPA Image Understanding Workshop*, Los Angeles, California, pp. 227-233 (23-25 February 1987). Also in *Proc. 6th National Conf. on Artificial Intelligence*, Seattle, Washington, pp. 706-711 (13-17 July 1987).

[Fukada 80a] Y. Fukada, "Spatial Clustering Procedures for Region Analysis," *Pattern Recognition*, Vol. 12, pp. 395-403 (1980).

[Fukada 80b] Y. Fukada, "Initial Segmentation and House Analysis in Color Aerophotographs," *Proc. 5th Int. Jnt. Conf. on Pattern Recognition*, Miami Beach, Florida, pp. 526-528 (December 1980).

[Fukushima 88] K. Fukushima, "A Neural Network for Visual Pattern Recognition," *IEEE Computer*, Vol. 21, No. 3, pp. 65-75 (March 1988).

[Garvey 76] T.D. Garvey, *Perceptual Strategies for Purposive Vision*, Technical Note 117, Artificial Intelligence Center, Stanford Research Institute, Menlo Park, California (September 1976).

[Glicksman 83] J. Glicksman, "Using Multiple Information Sources in a Computational Vision System," *8th Int. Jnt. Conf. on Artificial Intelligence*, Karlsruhe, West Germany, pp. 1078-1080 (8-12 August 1983).

[Goldberg 77] M. Goldberg and S. Shlien, "A Four-Dimensional Histogram Approach to the Clustering of LANDSAT Data," *Machine Processing of Remotely Sensed Data*, Purdue University, West Lafayette, Indiana, pp. 250-259 (21-23 June 1977).

[Goldberg 78] M. Goldberg and S. Shlien, "A Clustering Scheme for Multispectral Images," *IEEE Trans. on Systems, Man, and Cybernetics*, Vol. SMC-8, No. 2, pp. 86-92 (February 1978).

[Grinaker 80] S. Grinaker, "Edge Based Segmentation and Texture Separation," *Proc. 5th Int. Jnt. Conf. on Pattern Recognition*, Miami Beach, Florida, pp. 554-557 (1-4 December 1980).

[Gupta 74] J.N. Gupta and P.A. Wintz, "Computer Processing Algorithm for Locating Boundaries in Digital Pictures," *Proc. 2nd Int. Jnt. Conf. on Pattern Recognition*, Copenhagen, Denmark, pp. 155-156 (13-15 August 1974).

[Guzman 67] A. Guzman-Arenas, *Some Aspects of Pattern Recognition by Computer*, AI-TR-224 (MAC-TR-37), MIT AI Laboratory, Cambridge, Massachusetts (February 1967).

[Hanson 75] A.R. Hanson, E.M. Riseman, and P. Nagin, "Region Growing in Textured Outdoor Scenes," *Proc. 3rd Milwaukee Symp. on Automated Computation and Control*, Milwaukee, pp. 407-417 (April 1975). Also published as Report TR 75C-3, Dept. of Computer and Information Science, University of Massachusetts at Amherst, (February 1975).

[Hanson 80] A.R. Hanson and E.M. Riseman, "Processing Cones: A Computational Structure of Image Analysis," S. Tanimoto and A. Klinger (eds.), *Structured Computer Vision*, Academic Press, New York (1980). Also published as Computer and Information Science (COINS) Technical Report 81-38, University of Massachusetts at Amherst (December 1981).

[Haralick 69] R.M. Haralick and G.L. Kelly, "Pattern Recognition with Measurement Space and Spatial Clustering for Multiple Images," *Proc. IEEE*, Vol. 57, No. 4, pp. 654-665 (April 1969).

[Haralick 75] R.M. Haralick and I. Dinstein, "A Spatial Clustering Procedure for Multi-Image Data," *IEEE Trans. on Circuits and Systems*, Vol. CAS-22, No. 5, pp. 440-450 (May 1975).

[Harlow 79a] C.A. Harlow and R.W. Connors, "Image Segmentation and Texture Analysis," *Proc. SPIE Conf. on Image Understanding Systems II*, Vol. 205, San Diego, California, pp. 30-41 (1979).

[Harlow 79b] C.A. Harlow and R.W. Connors, "Image Analysis Segmentation Methods," in K.S. Fu and T. Pavlidis (eds.), *Biomedical Pattern Recognition and Processing*, Dahlem Konferenzen 1979, Berlin, Verlag Chemie GmbH, D-6940 Weinheim, pp. 111-130 (1979).

[Harlow 84] C.A. Harlow, R.W. Connors, and M.M. Trivedi, "A Computer Vision System for the Analysis of Aerial Scenes," *Proc. 7th Int. Jnt. Conf. on Pattern Recognition*, Montreal, Canada, pp. 407-410 (30 July - 2 August 1984).

[Harlow 85] C.A. Harlow, M. Trivedi, and R.W. Connors, "Use of Texture Operators in Segmentation," *Proc. SPIE Conf. on Applications of Artificial Intelligence*, Vol. 548, pp. 10-15+ (1985).

[Hartley 82] R.L. Hartley, L.J. Kitchen, C.-Y. Wang, and A. Rosenfeld, "Segmentation of FLIR Images: A Comparative Study," *Proc. DARPA Image Understanding Workshop*, Palo Alto, California, pp. 323-325 (15-16 September 1982).

[Harwood 87] D. Harwood, S. Chang, and L.S. Davis, "Interpreting Aerial Photographs by Segmentation and Search," *Proc. DARPA Image Understanding Workshop*, Los Angeles, California, pp. 507-520 (23-25 February 1987).

[Havens 83] W.S. Havens and A.K. Mackworth, "Representing Knowledge of the Visual World," *IEEE Computer*, Vol. 16, No. 10, pp. 90-96 (1983).

[Horowitz 74] S.L. Horowitz and T. Pavlidis, "Picture Segmentation by a Directed Split-and-Merge Procedure," *Proc. 2nd Int. Jnt. Conf. on Pattern Recognition*, Copenhagen, Denmark, pp. 424-433 (13-15 August 1974).

[Horowitz 76] S.L. Horowitz and T. Pavlidis, "Picture Segmentation by a Tree Traversal Algorithm," *J. of the ACM*, Vol. 23, No. 2, pp. 368-388 (April 1976).

[Hrechanyk 82] L.M. Hrechanyk and D.H. Ballard, "A Connectionist Model of Form Perception," *Proc. Workshop on Computer Vision: Representation and Control*, Rindge, New Hampshire, pp. 44-52 (23-25 August 1982). Also published in *Computer Science and Engineering Research Review*, Computer Science Department, University of Rochester (Fall 1982).

[Hrechanyk 83] L.M. Hrechanyk and D.H. Ballard, "Viewframes: A Connectionist Model of Form Perception," *Proc. DARPA Image Understanding Workshop*, Arlington, Virginia, pp. 64-76 (23 June 1983).

[Irwin 84] N. Irwin, *Combining Information from Independently Derived Intensity and Edge Segmentations*, Center for Automation Research, University of Maryland, College Park, Maryland (August 1984).

[Jarvis 73] R.A. Jarvis and E.A. Patrick, "Clustering Using a Similarity Measure Based on Shared Near Neighbors," *IEEE Trans. on Computers*, Vol. C-22, No. 11, pp. 1025-1034, (November 1973).

[Jarvis 75] R.A. Jarvis, "Image Segmentation by Interactively Combining Line, Region, and Semantic Structure," *Proc. Computer Graphics, Pattern Recognition, and Data Structures*, Los Angeles, California, pp. 279-288 (14-16 May 1975).

[Jarvis 77] R.A. Jarvis, "Region Based Image Segmentation Using Shared Near Neighbor Clustering," *Proc. 7th Int. Conf. on Cybernetics and Society*, Washington, D.C., pp. 641-647 (19-21 September 1972).

[Jayasimha 81] S. Jayasimha, *Some Techniques for the Automatic Interpretation of Aerial Photographs*, Report IPL-TR-81-009, Image Processing Laboratory, Rensselaer Polytechnic Institute, Troy, New York (April 1981).

[Kanade 77] T. Kanade, "Model Representation and Control Structures in Image Understanding," *Proc. 5th Int. Jnt. Conf. on Artificial Intelligence*, Cambridge, Massachusetts, pp. 1074-1082 (22-25 August 1977).

[Kanade 80] T. Kanade, "Region Segmentation: Signal vs Semantics," *Computer Graphics and Image Processing*, Vol. 13, pp. 279-297 (1980). Reprinted in [Pao 82], pp. 518-533.

[Kelly 70] M.D. Kelly, *Visual Identification of People by Computer*, Memo AI-130, STAN-CSD, Ph.D. Thesis (July 1970).

[Kender 76] J.R. Kender, *Saturation, Hue, and Normalized Color: Calculation, Digitization Effects, and Use*, Dept. of Computer Science, Carnegie-Mellon University, Pittsburgh, Pennsylvania (November 1976).

[Kender 77] J.R. Kender, "Instabilities in Color Transformations," *Proc. IEEE Conf. on Pattern Recognition and Image Processing*, Troy, New York, pp. 266-274 (6-8 June 1977).

[Keng 77] J. Keng, "Image Segmentation and Object Detection by a Syntactic Method," *Proc. DARPA Image Understanding Workshop*, Minneapolis, Minnesota, pp. 38-43 (April 1977).

[Kestner 80] W. Kestner, M. Bohner, R. Schärf, and M. Sties, "Object Guided Segmentation of Aerial Images," *Proc. 5th Int. Jnt. Conf. on Pattern Recognition*, Miami Beach, Florida, pp. 529-531 (December 1980). Reprinted in [Pao 82], pp. 557-559.

[Kestner 82] W. Kestner, "Segmentation and Abstract Interpretation in an Image Understanding System," *Proc. 6th Int. Jnt. Conf. on Pattern Recognition*, Munich, Germany, pp. 1011-1013 (19-22 October 1982).

[Khan 84b] N.A. Khan and R. Jain, "Matching an Imprecise Object Description with Models in a Knowledge Base," *Proc. 7th Int. Jnt. Conf. on Pattern Recognition*, Montreal, Canada, pp. 1131-1134 (30 July - 2 August 1984).

[Kohl 87] C.A. Kohl, A.R. Hanson, and E.M. Riseman, "Goal-Directed Control of Low-Level Process for Image Interpretation," *Proc. DARPA Image Understanding Workshop*, Los Angeles, California, pp. 538-551 (23-25 February 1987).

[Kohler 81] R. Kohler, "A Segmentation System Based on Thresholding," *Computer Graphics and Image Processing*, Vol. 15, pp. 319-338 (1981). See also *An Image Segmentation System based on Thresholding*, Report COINS-TR-78-23, University of Massachusetts at Amherst (December 1978).

[Laws 80a] K.I. Laws, *Textured Image Segmentation*, Ph.D. Thesis, Report USCIPI 940, Image Processing Institute, University of Southern California, Los Angeles (January 1980).

[Laws 80b] K.I. Laws, "Rapid Texture Identification," *Proc. SPIE Conf. on Image Processing for Missile Guidance*, San Diego, California, Vol. 238, pp. 376-380 (29 July - 1 August 1980).

[Laws 82] K.I. Laws, *The PHOENIX Image Segmentation System: Description and Evaluation*, Technical Note 289, Artificial Intelligence Center, SRI International, Menlo Park, California (December 1982).

[Laws 83] K.I. Laws, *The RELAX Image Relaxation System: Description and Evaluation*, Technical Note 301, Artificial Intelligence Center, SRI International, Menlo Park, California (August 1983).

[Laws 85] K.I. Laws, "Goal-Directed Textured-Image Segmentation," *Proc. SPIE Conf. on Applications of Artificial Intelligence II*, Vol. 548, Arlington, Virginia, (9-11 April 1985). Summarizes Technical Note 334, Artificial Intelligence Center, SRI International, Menlo Park, California (September 1984).

[Laws 88a] K.I. Laws, *Coarse Coding for Material and Object Identification* Technical Note 442, Artificial Intelligence Center, SRI International, Menlo Park, California (July 1988).

[Laws 88b] K.I. Laws, *Classification-Based Tracking of Objects and Materials* Technical Note 443, Artificial Intelligence Center, SRI International, Menlo Park, California (July 1988).

[Leclerc 88] Y.G. Leclerc, "Constructing Simple Stable Descriptions for Image Partitioning," *Proc. DARPA Image Understanding Workshop*, Cambridge, Massachusetts, pp. 365-382 (6-8 April 1988).

[Lee 83] H.Y. Lee, "Extraction of Textured Regions in Aerial Imagery," *Proc. DARPA Image Understanding Workshop*, Arlington, Virginia, pp. 298-303 (23 June 1983).

[Levine 76] M.D. Levine and J. Leemet, "A Method for Non-Purposive Picture Segmentation," *Proc. 3rd Int. Jnt. Conf. on Pattern Recognition*, Coronado, California, pp. 494-498 (8-11 November 1976).

[Levine 81a] M.D. Levine and D. Ting, "Intermediate Level Picture Interpretation Using Complete Two-Dimensional Models," *Computer Graphics and Image Processing*, Vol. 16, pp. 185-209 (1981).

[Levine 81b] M.D. Levine and S.I. Shaheen, "A Modular Computer Vision System for Picture Segmentation and Interpretation," *IEEE Trans. on Pattern Analysis and Machine Intelligence*, Vol. PAMI-3, No. 5, pp. 540-556 (September 1981). Also published in *Proc. IEEE Conf. on Pattern Recognition and Image Processing*, Chicago, Illinois, pp. 523-539 (August 1979).

[Lowe 85] D.G. Lowe, "Visual Recognition from Spatial Correspondence and Perceptual Organization," *Proc. 9th Int. Jnt. Conf. on Artificial Intelligence*, Los Angeles, pp. 953-959 (18-23 August 1985).

[Lumia 81] R. Lumia, R.M. Haralick, and L. Shapiro, "Texture Discrimination Using Region Based Primitives," *Proc. IEEE Conf. on Pattern Recognition and Image Processing*, Dallas, Texas, pp. 369-371 (3-5 August 1981).

[Lumia 83] R. Lumia, R.M. Haralick, O. Zuniga, L. Shapiro, T.-C. Pong, and F.-P. Wang, "Texture Analysis of Aerial Photographs," *Pattern Recognition*, Vol. 16, No. 1, pp. 39-46 (1983).

[Lutton 80] S.M. Lutton and O.R. Mitchell, "Adaptive Segmentation of Unique Objects," *Proc. 5th Int. Jnt. Conf. on Pattern Recognition*, Miami Beach, Florida, pp. 548-550 (December 1980).

[Mackworth 81] A.K. Mackworth and W.S. Havens, "Structuring Domain Knowledge for Visual Perception," *Proc. 7th Int. Jnt. Conf. on Artificial Intelligence*, Vancouver, British Columbia, Canada, pp. 625-627 (24-28 August 1981).

[Martelli 72] A. Martelli, "Edge Detection Using Heuristic Search Methods," *Computer Graphics and Image Processing*, Vol. 1, No. 2, pp. 169-182 (August 1972).

[Martelli 73] A. Martelli, "Contour Detection in Noisy Pictures Using Heuristic Search Methods," *Proc. 1st Int. Jnt. Conf. on Pattern Recognition*, Washington, D.C., pp. 375-388 (30 October-1 November 1973).

[Martelli 76] A. Martelli, "An Application of Heuristic Search Methods to Edge and Contour Detection," *Communications of the ACM*, Vol. 19, No. 2, pp. 73-83 (February 1976).

[Matsumoto 81] K. Matsumoto, M. Naka, and H. Yamanoto, "A New Clustering Method for LANDSAT Images Using Local Maximums of a Multidimensional Histogram," *Machine Processing of Remotely Sensed Data*, Purdue University, West Lafayette, Indiana, pp. 321-325 (23-26 June 1981).

[McKeown 84] D.M. McKeown, Jr., and J.L. Denlinger, "Map-Guided Feature Extraction from Aerial Imagery," *Proc. DARPA Image Understanding Workshop*, New Orleans, Louisiana, pp. 247-256 (3-4 October 1984). Also published as Report CMU-CS-84-117, Computer Science Department, Carnegie-Mellon University (3 April 1984).

[McKeown 85] D.M. McKeown, Jr., W.A. Harvey, Jr., and J. McDermott, "Rule-Based Interpretation of Aerial Imagery," *IEEE Trans. on Pattern Analysis and Machine Intelligence*, Vol. PAMI-7, No. 5, pp. 570-585 (September 1985). Also published in *Proc. IEEE Workshop on Principles of Knowledge-Based Systems*, pp. 145-157 (December 1984).

[Milgram 79] D.L. Milgram, "Region Extraction Using Convergent Evidence," *Computer Graphics and Image Processing*, Vol. 11, No. 1, pp. 1-12 (1979). Also published in *Proc. DARPA Image Understanding Workshop*, Minneapolis, Minnesota, pp. 58-64 (April 1977), and as Report TR-674, Computer Science Center, University of Maryland (June 1978).

[Milgram 81] D.L. Milgram and A. Rosenfeld, "Object Detection in Infrared Images," in L. Bolc and Z. Kulpa (eds.), *Digital Image Processing Systems*, Springer, New York, pp. 228-353 (1981).

[Minor 81] L.G. Minor and J. Sklansky, "The Detection and Segmentation of Blobs in Infrared Images," *IEEE Trans. on Systems, Man, and Cybernetics*, Vol. SMC-11, No. 3, pp. 194-201 (March 1981).

[Mitchell 78] O.R. Mitchell and S.G. Carlton, "Image Segmentation Using a Local Extrema Texture Measure," *Pattern Recognition*, Vol. 10, No. 3, pp. 205-210 (1978).

[Mitchell 79] O.R. Mitchell, S.P. Lutton, and S.P. Su, "Texture Image Segmentation Using Local Extrema," *Proc. IEEE Conf. on Pattern Recognition and Image Processing*, Chicago, Illinois, pp. 508-511 (August 1979).

[Montanari 71] U. Montanari, "On the Optimal Detection of Curves in Noisy Pictures," *Communications of the ACM*, Vol. 14, No. 5, pp. 335-345 (May 1971).

[Nagao 77] M. Nagao, T. Matsuyama, and H. Mori, "Structured Analysis of Complex Aerial Photographs," *Proc. 5th Int. Jnt. Conf. on Artificial Intelligence*, Cambridge, Massachusetts, pp. 610-616 (22-25 August 1977).

[Nagao 79] M. Nagao, T. Matsuyama, and Y. Ikeda, "Region Extraction and Shape Analysis in Aerial Photographs," *Computer Graphics and Image Processing*, Vol. 10, No. 3, pp. 195-223 (July 1979). Also published in *Proc. 4th Int. Jnt. Conf. on Pattern Recognition*, Kyoto, Japan, pp. 620-628 (November 1978).

[Nagao 80] M. Nagao and T. Matsuyama, *A Structural Analysis of Complex Aerial Photographs*, Plenum Press, New York, (1980).

[Nagao 84] M. Nagao, "Control Strategies in Pattern Analysis," *Pattern Recognition*, Vol. 17, No. 1, pp. 45–56 (1984). Also published in *Proc. 6th Int. Jnt. Conf. on Pattern Recognition*, Munich, Germany, pp. 996–1006 (19–22 October 1982).

[Nagin 77] P.A. Nagin, A.R. Hanson, and E.M. Riseman, *Region Extraction and Description Through Planning*, Computer and Information Science Report 77–8, University of Massachusetts at Amherst (May 1977).

[Nagin 78] P.A. Nagin, *Segmentation Using Spatial Context and Feature Space Cluster Labels*, Computer and Information Science Report 78–8, University of Massachusetts at Amherst (May 1978).

[Nagin 79] P. Nagin, *Studies in Image Segmentation Algorithms based on Histogram Clustering and Relaxation*, Ph.D. Thesis, COINS Technical Report 79–15, University of Massachusetts at Amherst (September 1979).

[Nagin 82] P. Nagin, A. Hanson, and E. Riseman, "Studies in Global and Local Histogram-Guided Relaxation Algorithms," *IEEE Trans. on Pattern Analysis and Machine Intelligence*, Vol. PAMI-4, No. 3, pp. 263–277 (May 1982).

[Nagy 72] G. Nagy and J. Tolaba, "Nonsupervised Crop Classification Through Airborne Multispectral Observation," *IBM Journal of Research and Development*, Vol. 16, No. 2, pp. 138–153 (March 1972).

[Nahi 77] N.E. Nahi and M.H. Jahanshahi, "Image Boundary Estimation," *IEEE Trans. on Computers*, Vol. C-26, No. 8, pp. 772–781 (August 1977). Reprinted in H.C. Andrews (ed.), *Tutorial and Selected Papers in Digital Image Processing*, IEEE Computer Society, Long Beach, California, pp. 546–555 (1978).

[Nahi 78] N.E. Nahi and S. Lopez-Mora, "Estimation-Detection of Object Boundaries in Noisy Images," *IEEE Trans. on Automatic Control*, Vol. AC-23, No. 5, pp. 834–846 (October 1978).

[Nakagawa 78a] Y. Nakagawa and A. Rosenfeld, *Some Experiments on Variable Thresholding*, Report TR-626, University of Maryland, College Park, Maryland (January 1978).

[Nakagawa 78b] Y. Nakagawa and A. Rosenfeld, "A Note on the Use of Local Min and Max Operations in Digital Picture Processing," *IEEE Trans. on Systems, Man and Cybernetics*, Vol. SMC-8, No. 8, pp. 632–635 (August 1978).

[Narayanan 81] K.A. Narayanan and A. Rosenfeld, "Image Smoothing by Local Use of Global Information," *IEEE Tran. on Systems, Man, and Cybernetics*, pp. 826–831 (1981).

[Narendra 77] P.M. Narendra and M. Goldberg, "A Non-Parametric Clustering Scheme for LANDSAT," *Pattern Recognition*, Vol. 9, pp. 207–215 (1977).

[Nazif 84] A.M. Nazif and M.D. Levine, "Low Level Image Segmentation: An Expert System," *IEEE Trans. on Pattern Analysis and Machine Intelligence*, Vol. PAMI-6, No. 5, pp. 555–577 (September 1984). Also published as Technical Report TR-83-4, Dept. of Electrical Engineering, McGill University (April 1983).

[Nevatia 76] R. Nevatia, "Locating Object Boundaries in Textured Environments," *IEEE Trans. on Computers*, Vol. C-25, No. 11, pp. 1170–1175 (November 1976).

[Nevatia 80] R. Nevatia and K.R. Babu, "Linear Feature Extraction and Description," *Computer Graphics and Image Processing*, Vol. 13, pp. 257–269 (1980); reprinted in R. Chellappa and A.A. Sawchuk (eds.), *Digital Image Processing and Analysis: Volume 2: Digital Image Analysis*, IEEE Computer Society Press, Washington, D.C., pp. 81–93 (1985). Also published in *Proc. Int. Jnt. Conf. on Artificial Intelligence*, pp. 639–641 (1979).

[O'Gorman 73] F. O'Gorman and M.B. Clowes, "Finding Picture Edges Through Collinearity of Feature Points," *Proc. 3rd Int. Jnt. Conf. on Artificial Intelligence*, Stanford, California, pp. 543–555, (20–23 August 1973).

[Ohlander 75] R.B. Ohlander, *Analysis of Natural Scenes*, PhD. Thesis, Dept. of Computer Science, Carnegie-Mellon University, Pittsburgh, Pennsylvania (April 1975).

[Ohlander 78] R. Ohlander, K. Price, and D.R. Reddy, "Picture Segmentation Using a Recursive Region Splitting Method," *Computer Graphics and Image Processing*, Vol. 8, No. 3, pp. 313-333 (December 1978).

[Ohta 80] Y. Ohta, *A Region-Oriented Image-Analysis System by Computer*, Ph.D. Thesis, Dept. of Information Sciences, Kyoto University, Kyoto, Japan (March 1980).

[Otsu 79] N. Otsu, "A Threshold Selection Method from Gray-Level Histograms," *IEEE Trans. on Systems, Man, and Cybernetics*, Vol. SMC-9, No. 1, pp. 62-66 (January 1979).

[Panda 77] D.P. Panda, "Segmentation of FLIR Images by Pixel Classification," *Proc. DARPA Image Understanding Workshop*, Minneapolis, Minnesota, pp. 65-70 (April 1977). Also published as Technical Report TR-508, Computer Science Center, University of Maryland, College Park, Maryland (February 1977).

[Panda 78] D.P. Panda and A. Rosenfeld, "Image Segmentation by Pixel Classification in (Grey Level, Edge Value) Space," *IEEE Trans. on Computers*, Vol. C-27, No. 9, pp. 875-879 (September 1978).

[Pao 82] Y.-H. Pao and G.W. Ernst (eds.), *Context-Directed Pattern Recognition and Machine Intelligence Techniques for Information Processing*, IEEE Computer Society Press, Silver Spring, Maryland (1982).

[Parvin 83] B.A. Parvin and B. Bhanu, "Segmentation of Images Using a Relaxation Technique," *Proc. IEEE Conf. on Computer Vision and Pattern Recognition*, Washington, D.C., pp. 151-153 (19-23 June 1983).

[Pavlidis 72] T. Pavlidis, "Segmentation of Pictures and Maps through Functional Approximation," *Computer Graphics and Image Processing*, Vol. 1, pp. 360-372 (1972).

[Pavlidis 75] T. Pavlidis and S. Tanimoto, "Texture Identification by a Directed Split-and-Merge Procedure," *Proc. Computer Graphics, Pattern Recognition, and Data Structures*, Los Angeles, California, pp. 201-203 (14-16 May 1975).

[Pentland 84] A.P. Pentland, "Fractal-Based Description of Natural Scenes," *IEEE Trans. on Pattern Analysis and Machine Intelligence*, Vol. PAMI-6, No. 6, pp. 661-674 (November 1984), reprinted in R. Chellappa and A..A. Sawchuk (eds.), *Digital Image Processing and Analysis: Volume 2: Digital Image Analysis*, pp. 368-381 (1985), and in A.P. Pentland (ed.), *From Pixels to Predicates: Recent Advances in Computational and Robotic Vision*, Ablex, Norwood, New Jersey pp. 227-252 (1985). Also published in *Proc. IEEE Conf. on Computer Vision and Pattern Recognition*, Washington, D.C., pp. 201-209 (19-23 June 1983), and in *Proc. DARPA Image Understanding Workshop*, Arlington, Virginia, pp. 184-192 (23 June 1983).

[Pentland 86a] A.P. Pentland, "Perceptual Organization and the Representation of Natural Form," *Artificial Intelligence*, Vol. 28, No. 3, pp. 293-331 (May 1986). Also published as Technical Note 357, Artificial Intelligence Center, SRI International, Menlo Park, California (July 1985, revised July 1986).

[Pentland 86b] A.P. Pentland, "Shading into Texture," *Artificial Intelligence*, Vol. 29, No. 2, pp. 147-170 (August 1986). Also published in *Proc. 4th National Conf. on Artificial Intelligence*, Austin, Texas, pp. 269-273 (6-10 August 1984), in *Proc. DARPA Image Understanding Workshop*, New Orleans, Louisiana, pp. 179-183 (3-4 October 1984), and in [Pentland 85a], pp. 253-267 (1985).

[Pietikäinen 81] M. Pietikäinen and A. Rosenfeld, "Image Segmentation by Texture Using Pyramid Node Linking," *IEEE Trans. on Systems, Man, and Cybernetics*, Vol. SMC-11, No. 12, pp. 822-825 (December 1981).

[Pietikäinen 82a] M. Pietikäinen, *Image Texture Analysis and Segmentation*, Ph.D. Thesis, Department of Electrical Engineering, University of Oulu, Oulu, Finland, Acta Universitatis Ouluensis, Series C, Technica No. 21, Electronica No. 1 (February 1982).

[Pietikäinen 82b] M. Pietikäinen, A. Rosenfeld, and I. Walter, "Split-and-Link Algorithms for Image Segmentation," *Pattern Recognition*, Vol. 15, No. 4, pp. 287-298 (1982).

[Pingle 71] K. Pingle and J. Tenenbaum, "An Accommodating Edge Follower," *Proc. 2nd Int. Jnt. Conf. on Artificial Intelligence*, London, pp. 1-7 (1-3 September 1971).

[Pong 81] T.C. Pong, L.G. Shapiro, and R.M. Haralick, "A Facet Model Region Growing Algorithm," *Proc. IEEE Conf. on Pattern Recognition and Image Processing*, Dallas, Texas, pp. 279-284 (August 1981).

[Pong 84] T.C. Pong, L.G. Shapiro, L.T. Watson, and R.M. Haralick, "Experiments in Segmentation Using a Facet Model Region Grower," *Computer Vision, Graphics, and Image Processing* Vol. 25, (January 1984).

[Prewitt 66] J.M.S. Prewitt and M.L. Mendelson, "The Analysis of Cell Images," *Ann. N.Y. Acad. of Sci.*, Vol. 128, pp. 1035-1053 (1966).

[Prewitt 70] J.M.S. Prewitt, "Object Enhancement and Extraction," in B.S. Lipkin and A. Rosenfeld (eds.), *Picture Processing and Psychopictorics*, Academic Press, New York, pp. 75-149 (1970).

[Price 76] K.E. Price, *Change Detection and Analysis in Multi-Spectral Images*, Ph.D. Thesis, Dept. of Computer Science, Carnegie-Mellon University, Pittsburgh, Pennsylvania (December 1976).

[Price 81] K.E. Price, "Relaxation Matching Applied to Aerial Images," *Proc. DARPA Image Understanding Workshop*, Washington, D.C., pp. 22-25 (23 April 1981).

[Quam 78] L.H. Quam, "Road Tracking and Anomaly Detection in Aerial Imagery," *Proc. DARPA Image Understanding Workshop*, Cambridge, Massachusetts, pp. 51-55 (3-4 May 1978). Also published as Technical Note 158, Artificial Intelligence Center, SRI International, Menlo Park, California (March 1978).

[Raafat 80] H.M. Raafat and A.K.C. Wong, "Texture Information Directed Algorithm for Biomedical Image Segmentation and Classification," *Proc. Conf. on Cybernetics and Society*, pp. 1003-1008 (1980).

[Raafat 86] H.M. Raafat and A.K.C. Wong, "Texture-Based Image Segmentation," *Proc. IEEE Conf. on Computer Vision and Pattern Recognition*, Miami Beach, Florida, pp. 469-475 (22-26 June 1986).

[Ramer 75] U. Ramer, "Extraction of Line Structures from Photographs of Curved Objects," *Computer Graphics and Image Processing*, Vol. 4, pp. 81-103 (June 1975).

[Rassbach 80] M.E. Rassbach, "CLASSY: An Adaptive Clustering Algorithm," *Proc. 5th Int. Jnt. Conf. on Pattern Recognition*, Miami Beach, Florida, pp. 442-444 (December 1980).

[Reynolds 84] G. Reynolds, N. Irwin, A. Hanson, and E. Riseman, "Hierarchical Knowledge-Directed Object Extraction Using a Combined Region and Line Representation," *Proc. DARPA Image Understanding Workshop*, New Orleans, Louisiana, pp. 195-204 (3-4 October 1984). Also published in *Proc. Workshop on Computer Vision: Representation and Control*, Annapolis, Maryland, pp. 238-247 (30 April - 2 May 1984).

[Ridler 78] T.W. Ridler and S. Calvard, "Picture Thresholding Using an Iterative Selection Method," *IEEE Trans. on Systems, Man, and Cybernetics*, Vol. SMC-8, No. 8, pp. 630-632 (August 1978).

[Riseman 77] E.M. Riseman and M.A. Arbib, "Computational Techniques in the Visual Segmentation of Static Scenes," *Computer Graphics and Image Processing*, Vol. 6, No. 3, pp. 221-276 (June 1977).

[Robertson 73a] T.V. Robertson, K.S. Fu, and P.H. Swain, *Multispectral Image Partitioning*, LARS Information Note 071373, Technical Report TR-EE 73-26, School of Electrical Engineering, Purdue University, West Lafayette, Indiana (August 1973).

[Robertson 73b] T.V. Robertson, "Extraction and Classification of Objects in Multispectral Images," *Machine Processing of Remotely Sensed Data*, Purdue University, West Lafayette, Indiana, pp. 3B-27-3B-34 (16-18 October 1973).

[Rosenfeld 78] A. Rosenfeld and L.S. Davis, "Iterative Histogram Modification," *IEEE Trans. on Systems, Man, and Cybernetics*, Vol. SMC-8, No. 4, pp. 300-302 (April 1978). Also published as TR-519, Computer Science Center, University of Maryland (April 1977).

[Rosenfeld 79a] A. Rosenfeld, A. Danker, and C.R. Dyer, "Blob Extraction by Relaxation," *Proc. DARPA Image Understanding Workshop*, Palo Alto, California, pp. 61-65 (24-25 April 1979).

[Rosenfeld 79b] A. Rosenfeld and L.S. Davis, "Image Segmentation and Image Models," *Proc. IEEE*, Vol. 67, No. 5, pp. 764-772 (May 1979).

[Rosenthal 78] D.A. Rosenthal, *An Inquiry Driven Computer Vision System based on Visual and Conceptual Hierarchies*, Ph.D. Thesis, University of Pennsylvania, Philadelphia (1978).

[Rosenthal 84] D.A. Rosenthal and R. Bajcsy, "Visual and Conceptual Hierarchy: A Paradigm for Studies of Automated Generation of Recognition Strategies," *IEEE Trans. on Pattern Analysis and Machine Intelligence*, Vol. PAMI-6, No. 3, pp. 319-325 (May 1984).

[Rubin 78] S.M. Rubin, *The ARGOS Image Understanding System*, Ph.D. Thesis, Carnegie-Mellon University (1978). Summarized in *Proc. DARPA Image Understanding Workshop*, Pittsburgh, Pennsylvania, pp. 159-162 (14-15 November 1978).

[Rubin 80] S.M. Rubin, "Natural Scene Recognition Using LOCUS Search," *Computer Graphics and Image Processing*, Vol. 13, No. 4, pp. 298-333 (1980).

[Russell 79] D.M. Russell, "Where Do I Look Now?: Modeling and Inferring Object Locations by Constraints," *Proc. IEEE Conf. on Pattern Recognition and Image Processing*, Chicago, Illinois, pp. 175-183 (August 1979).

[Saba 78] T.G. Saba, D.J. Bryant, and E.L. Hall, "Scene Segmentation Using Uniform Clustering," *Proc. SOUTHEASTCON '78 Region 3 Conference*, Atlanta, pp. 217-221 (April 1978).

[Sabbah 82] D. Sabbah, *A Connectionist Approach to Visual Recognition*, Ph.D. Thesis, TR107, Computer Science Department, University of Rochester (April 1982).

[Sadjadi 79] F.A. Sadjadi and E.L. Hall, "Segmentation Using Material Information," *Proc. SPIE Conf. on Digital Processing of Aerial Images*, Vol. 186, pp. 254-259 (1979).

[Sakai 76] T. Sakai, T. Kanade, and Y. Ohta, "Model Based Interpretation of Outdoor Scenes," *Proc. 3rd Int. Jnt. Conf. on Pattern Recognition*, Coronado, California, pp. 581-585 (8-11 November 1976).

[Sarabi 81] A. Sarabi and J.K. Aggarwal, "Segmentation of Chromatic Images," *Pattern Recognition*, Vol. 13, No. 6, pp. 417-427 (1981).

[Schachter 75] B.J. Schachter, L.S. Davis, and A. Rosenfeld, *Scene Segmentation by Cluster Detection in Color Spaces*, Report TR-424, Computer Science Center, University of Maryland (November 1975). Also in *SIGART*, No. 58, pp. 16-17 (June 1976).

[Schachter 77] B.J. Schachter, L.S. Davis, and A. Rosenfeld, *Some Experiments in Image Segmentation by Clustering of Local Feature Values*, Report TR-510, Computer Science Center, University of Maryland (March 1977).

[Schachter 79] B.J. Schachter, L.S. Davis, and A. Rosenfeld, "Some Experiments in Image Segmentation by Clustering of Local Feature Values," *Pattern Recognition*, Vol. 11, No. 1, pp. 19-28 (1979).

[Seibert 88] M. Seibert and A.M. Waxman, "Early-Vision Applications of Feature-Map Diffusion-Enhancement Networks," *Proc. AAAI 1988 Spring Symposium Series, Physical and Biological Approaches to Computational Vision*, pp. 124-125 (22-24 March 1988).

[Selfridge 81] P.G. Selfridge, "Reasoning About Images: Using Meta-Knowledge in Aerial Image Understanding," *Proc. 7th Int. Jnt. Conf. on Artificial Intelligence*, Vancouver, British Columbia, Canada, pp. 755-757 (24-28 August 1981).

[Selfridge 82] P.G. Selfridge, *Reasoning about Success and Failure in Aerial Image Understanding*, Ph.D. Thesis, TR-103, Computer Science Dept., University of Rochester, New York (May 1982). Summarized in P.G. Selfridge and K.R. Sloan, *Proc. IEEE Conf. on Pattern Recognition and Image Processing*, Las Vegas, Nevada, pp. 44-49 (14-17 June 1982).

[Shafer 80] S.A. Shafer, *MOOSE: User's Manual, Implementation Guide, Evaluation*, IFI-HH B-70/80, Fachbereich Informatik, University of Hamburg, W. Germany (April 1980).

[Shafer 82] S. Shafer and T. Kanade, "Recursive Region Segmentation by Analysis of Histograms," *Proc. Int. Conf. on Acoustics, Speech, and Signal Processing*, Paris, France, pp. 1166-1171 (May 1982).

[Shaheen 79] S.I. Shaheen, *Image Segmentation and Interpretation Using a Knowledge Data Base*, Ph.D. Thesis, Dep. Elec. Eng., McGill University, Montreal (1979). Summarized in the *2nd Nat. Conf. of the Canadian Society for Computational Studies of Intelligence*, Toronto, Canada.

[Shirai 72] Y. Shirai, "A Heterarchical Program for Recognition of Polyhedra," *Artificial Intelligence*, Vol. 4, No. 2 (Summer 1973). Also in P.H. Winston (ed.), *The Psychology of Computer Vision*, McGraw-Hill, New York (1975), and published as Artificial Intelligence Memo 263, Massachusetts Institute of Technology (June 1972).

[Silverman 86] J.F. Silverman and D.B. Cooper, *Bayesian Clustering for Unsupervised Estimation of Surface and Texture Models*, Technical Report LEMS-20, Laboratory for Engineering Man/Machine Systems, Brown University (April 1986).

[Silverman 87] J.F. Silverman and D.B. Cooper, "Unsupervised Bayesian Model-Learning with Applications to Textured and Polynomial Image Segmentation," *Proc. 1st Int. Conf. on Computer Vision*, London, England, pp. 672-676, (8-11 June 1987).

[Skevington 81] C. Skevington, G.M. Flachs, and B. Schaming, "A Statistical Approach to Image Segmentation," *Proc. IEEE Conf. on Pattern Recognition and Image Processing*, Dallas, Texas, pp. 267-272 (August 1981).

[Sloan 77] K.R. Sloan, Jr., *World Model Driven Recognition of Natural Scenes*, Ph.D. Thesis, The Moore School of Electrical Engineering, University of Pennsylvania, Philadelphia (June 1977).

[Sloan 79] K.R. Sloan and R. Bajcsy, *World Model Driven Recognition of Outdoor Scenes*, TR40, Computer Science Department, University of Rochester (September 1979).

[Somerville 76] C. Somerville and J.L. Mundy, "One Pass Contouring of Images Through Planar Approximation," *Proc. 3rd Int. Conf. on Pattern Recognition*, Coronado, California, pp. 745-748 (November 1976).

[Strat 86] T.M. Strat and M.A. Fischler, "One-Eyed Stereo: A General Approach to Modeling 3-D Scene Geometry," *IEEE Trans. on Pattern Analysis and Machine Intelligence*, Vol. PAMI-8, No. 6, pp. 730-741 (November 1986). Also published in *Proc. 9th Int. Int. Conf. on Artificial Intelligence*, Los Angeles, California, pp. 937-943 (18-23 August 1985), in *Proc. DARPA Image Understanding Workshop*, Miami Beach, Florida, pp. 363-372 (9-10 December 1985), and as *One-Eyed Stereo: A Unified Strategy to Recover Shape from a Single Image*, Technical Note 367, Artificial Intelligence Center, SRI International, Menlo Park, California (5 November 1985).

[Swain 68] P.H. Swain and K.S. Fu, "On the Applications of Nonparametric Techniques to Crop Classification Problems," *National Electronics Conf. Proc.*, Vol. 24, pp. 14-19 (1968).

[Swain 81] P.H. Swain, S.B. Vardeman and J.C. Tilton, "Contextual Classification of Multispectral Image Data," *Pattern Recognition*, Vol. 13, No. 6, pp. 429-441 (1981).

[Suk 83] M. Suk and S.-M. Chung, "A New Image Segmentation Technique Based on Partition Mode Test," *Pattern Recognition*, Vol. 16, No. 5, pp. 469-480 (1983).

[Sze 82] T.-W. Sze and Y.-H. Yang, "Goal Directed Segmentation," *Proc. IEEE Conf. on Pattern Recognition and Image Processing*, Las Vegas, Nevada, pp. 504-511 (14-17 June 1982).

[Tanimoto 76] S.L. Tanimoto and T. Pavlidis, "Graph Labelling Algorithms for Picture Analysis," *Proc. 3rd Int. Int. Conf. on Pattern Recognition*, Coronado, California, pp. 749-752 (8-11 November 1976).

[Tanimoto 77] S.L. Tanimoto and T. Pavlidis, "The Editing of Picture Segmentations Using Local Analysis of Graphs," *Communications of the ACM*, Vol. 20, No. 4, pp. 223-229 (April 1977).

[Tanimoto 82] S.L. Tanimoto, "Paradigms for Control of Vision Using Inference in Networks," *Proc. Workshop on Computer Vision: Representation and Control*, Rindge, New Hampshire, pp. 3-13 (23-25 August 1982).

[Taylor 76] R.H. Taylor, *Generating AI Programs from High-Level Task Descriptions*, Ph.D. Thesis, Stanford AI Lab (1976).

[Tenenbaum 76] J.M. Tenenbaum and H.G. Barrow, "IGS: A Paradigm for Integrating Image Segmentation and Interpretation," *Proc. 3rd Int. Int. Conf. on Pattern Recognition*, Coronado, California, pp. 504-513 (8-11 November 1976).

[Tenenbaum 77] J.M. Tenenbaum and H.G. Barrow, "Experiments in Interpretation Guided Segmentation," *Artificial Intelligence*, Vol. 8, No. 3, pp. 241-274 (June 1977). Also published as Technical Note 123, Artificial Intelligence Center, Stanford Research Institute, Menlo Park, California (March 1976).

[Thompson 74] W.B. Thompson, *The Role of Texture in Computerized Scene Analysis*, USCIP Report 550, Image Processing Institute, University of Southern California, Los Angeles (December 1974).

[Tomita 73] F. Tomita, M. Yachida, and S. Tsuji, "Detection of Homogeneous Regions by Structural Analysis," *Proc. 3rd Int. Int. Conf. on Artificial Intelligence*, Stanford, California, pp. 564-571 (20-23 August 1973).

[Tomita 77] F. Tomita and S. Tsuji, "Extraction of Multiple Regions by Smoothing in Selected Neighborhoods," *IEEE Trans. on Systems, Man, and Cybernetics*, Vol. SMC-7, No. 2, pp. 107-109 (February 1977).

[Triendl 74] E.E. Triendl, "Recognition of Texture Edges by Propagation," *Proc. 2nd Int. Jnt. Conf. on Pattern Recognition*, Copenhagen, Denmark, pp. 284–285 (13–15 August 1974).

[Trivedi 84a] M.M. Trivedi, C.A. Harlow, R.W. Conners, and S. Goh, "Object Detection Based on Gray Level Cooccurrence," *Computer Vision, Graphics, and Image Processing*, Vol. 28, pp. 199–219 (1984).

[Trivedi 84b] M.M. Trivedi, C.A. Harlow, and S. Goh, "An Optimal Approach for the Incorporation of Context," Technical Report, Remote Sensing and Image Processing Laboratory, Dept. of Electrical Engineering, Louisiana State University, Baton Rouge, Louisiana (September 1984).

[Trivedi 84c] M.M. Trivedi and J.C. Bezdek, *Low-Level Segmentation of Aerial Images with Fuzzy Clustering*, Research Study in Electrical and Computer Engineering, Louisiana State University, Baton Rouge, Louisiana (October 1984).

[Trivedi 85] M.M. Trivedi, "Analysis of Aerial Images Using Fuzzy Clustering," in J.C. Bezdek (ed.), *The Analysis of Fuzzy Information*, CRC Press (1985).

[Tsai 78] W.H. Tsai and K.S. Fu, "Image Segmentation and Recognition by Texture Discrimination: A Syntactic Approach," *Proc. 4th Int. Jnt. Conf. on Pattern Recognition*, Kyoto, Japan (November 1978).

[Tsuji 73] S. Tsuji and F. Tomita, "A Structural Analyzer for a Class of Textures," *Computer Graphics and Image Processing*, Vol. 2, No. 3/4, pp. 216–231 (December 1973).

[Uhr 72] L. Uhr, "Layered 'Recognition Cone' Networks that Preprocess, Classify and Describe," *IEEE Trans. on Computers*, Vol. C-21, No. 7, pp. 758–768 (July 1972).

[Uhr 82] L. Uhr and L. Schmitt, *The Several Steps from ICON to SYMBOL, Using Structured Cone/Pyramids*, Technical Report 481 Wisconsin Univ. Computer Science Dept., (1982).

[Walters 86] D. Walters, "Selection and Use of Image Features for Segmentation of Boundary Images," *Proc. IEEE Conf. on Computer Vision and Pattern Recognition*, Miami Beach, Florida, pp. 319–324 (22–26 June 1986).

[Wechsler 75] H. Wechsler and J. Sklansky, "Automatic Detection of Rib Contours in Chest Radiographs," *Proc. 4th Int. Jnt. Conf. on Artificial Intelligence*, Tbilisi, Georgia, USSR (5–6 September 1975).

[Wechsler 77] H. Wechsler and J. Sklansky, "Finding the Rib Cage in Chest Radiographs," *Pattern Recognition*, Vol. 9, No. 1, pp. 21–30 (January 1977).

[Wechsler 79] H. Wechsler and M. Kidode, "A Random Walk Procedure for Texture Discrimination," *IEEE Trans. on Pattern Analysis and Machine Intelligence*, Vol. PAMI-1, No. 3, pp. 272–280 (1979).

[Wechsler 80] H. Wechsler, "Taxonomy and Segmentation of Textured Images," *Proc. 5th Int. Conf. on Pattern Recognition*, Miami, Florida, pp. 532–534 (1–4 December 1980).

[Wesley 82a] L.P. Wesley and A.R. Hanson, "The Use of an Evidential-Based Model for Representing Knowledge and Reasoning about Images in the VISIONS System," *Proc. IEEE Workshop on Computer Vision: Representation and Control*, Rindge, New Hampshire, pp. 14–25 (23–25 August 1982).

[Wesley 82b] L. Wesley, *Control as an Inference Process: A Model for Making Resource Allocation Decisions in the VISIONS System*, COINS Technical Report, University of Massachusetts (1982).

[Wesley 86] L.P. Wesley, "Evidential Knowledge-Based Computer Vision," *Optical Engineering*, Vol. 25, No. 3, pp. 363-379 (March 1986). Also published as Technical Note 374, Artificial Intelligence Center, SRI International, Menlo Park, California (January 1986).

[Weszka 78a] J.S. Weszka, "A Survey of Threshold Selection Techniques," *Computer Graphics and Image Processing*, Vol. 7, No. 2, pp. 259-265 (April 1978).

[Weszka 78b] J.S. Weszka and A. Rosenfeld, "Threshold Evaluation Techniques," *IEEE Trans. on Systems, Man, and Cybernetics*, Vol. SMC-8, No. 8, pp. 622-629 (August 1978).

[Weszka 79] J.S. Weszka and A. Rosenfeld, "Histogram Modification for Threshold Selection," *IEEE Trans. on Systems, Man, and Cybernetics*, Vol. SMC-9, No. 1, pp. 38-52 (January 1979).

[Weymouth 81] T.E. Weymouth, "Experiments in Knowledge-Driven Interpretation of Natural Scenes," *Proc. 7th Int. Jnt. Conf. on Artificial Intelligence*, Vancouver, British Columbia, Canada, pp. 628-630 (24-28 August 1981).

[Weymouth 83] T.E. Weymouth, J.S. Griffith, A.R. Hanson, and E.M. Riseman, "Rule Based Strategies for Image Interpretation," *Proc. National Conf. on Artificial Intelligence*, Washington, D.C., pp. 429-432 (22-26 August 1983). A longer version was published in *Proc. DARPA Image Understanding Workshop*, Arlington, Virginia, pp. 193-202 (23 June 1983).

[Wharton 83] S.W. Wharton, "A Generalized Histogram Clustering Scheme for Multidimensional Image Data," *Pattern Recognition*, Vol. 16, No. 2, pp. 193-199 (1983).

[Witkin 80] A.P. Witkin, "A Statistical Technique for Recovering Surface Orientation from Texture in Natural Imagery," *Proc. 1st Annual Nat. Conf. on Artificial Intelligence*, Stanford, California, pp. 1-3 (August 1980).

[Witkin 81a] A.P. Witkin, "Recovering Surface Shape and Orientation from Texture," *Artificial Intelligence*, Vol. 17, No. 1-3, pp. 17-45 (1981).

[Witkin 81b] A.P. Witkin, *Recovering Intrinsic Scene Characteristics from Images*, Interim Report, SRI Project 1019, Artificial Intelligence Center, SRI International, Menlo Park, California (September 1981).

[Witkin 83] A.P. Witkin, "Scale-Space Filtering," *Proc. 8th Int. Jnt. Conf. on Artificial Intelligence*, Karlsruhe, West Germany, pp. 1019-1022 (8-12 August 1983).

[Yachida 71] M. Yachida and S. Tsuji, "Application of Color Information to Visual Perception," *Pattern Recognition*, Vol. 3, No. 3, pp. 307-323 (October 1971).

[Yakimovsky 73a] Y. Yakimovsky, *Scene Analysis Using a Semantic Base for Region Growing*, Ph.D. Thesis, AIM-209, STAN-CS-73-380, Computer Science Dept., Stanford University (June 1973).

[Yakimovsky 73b] Y. Yakimovsky and J.A. Feldman, "A Semantics-Based Decision Theory Region Analyzer," *Proc. 3rd Int. Jnt. Conf. on Artificial Intelligence*, Stanford, California, pp. 580-588 (20-23 August 1973). Reprinted in J.K. Aggarwal, R.O. Duda, and A. Rosenfeld (eds.), *Computer Methods in Image Processing*, IEEE Press, New York, pp. 426-434 (1977).

[Yakimovsky 74] Y. Yakimovsky, *On the Recognition of Complex Structures: Computer Software Using Artificial Intelligence Applied to Pattern Recognition*, Technical Memorandum 33-693, Jet Propulsion Laboratory, Pasadena, California (May 1974).

[Yakimovsky 76] Y. Yakimovsky, "Boundary and Object Detection in Real World Images," *J. of the ACM*, Vol. 23, No. 4, pp. 599-618 (October 1976). Also published as Technical Memorandum 33-709, Jet Propulsion Laboratory, Pasadena, California (November 1974).

[Yoo 78] J.R. Yoo and T.S. Huang, *Image Segmentation by Unsupervised Clustering and its Applications*, Report TR-EE-78-19, Purdue Univ., West Lafayette, Indiana (1978).

- [Zucker 75] S.W. Zucker, A. Rosenfeld, and L.S. Davis, "Picture Segmentation by Texture Discrimination," *IEEE Trans. on Computers*, Vol. C-24, No. 12, pp. 1228-1233 (December 1975).
- [Zucker 76] S.W. Zucker, "Region Growing: Childhood and Adolescence," *Computer Graphics and Image Processing*, Vol. 5, No. 3, pp. 382-399 (September 1976).